

Host: subir



Amit

## Agenda

### Day 1 : Current state


<a href="#">Keynote by Mike</a>
<a href="#">AdSpam overview</a>
break
<a href="#">Planning process</a>
<a href="#">Future of AdSpam Infra</a>
break
<a href="#">Metrics</a>
break
<a href="#">Privacy</a>


### Day 2: Deep dives into key challenges

Keynote by Saikat
High level strategy overview
break
Trust Graph powered analysis
break
Challenges with supporting new/emerging markets/ segments/ products
Improving precision and monitoring
break
Impression Spam
Apps & Play
Closing Remarks
Fun event

See detailed agenda at [go/adspam-strategysummit20-agenda](https://go/adspam-strategysummit20-agenda)

Google






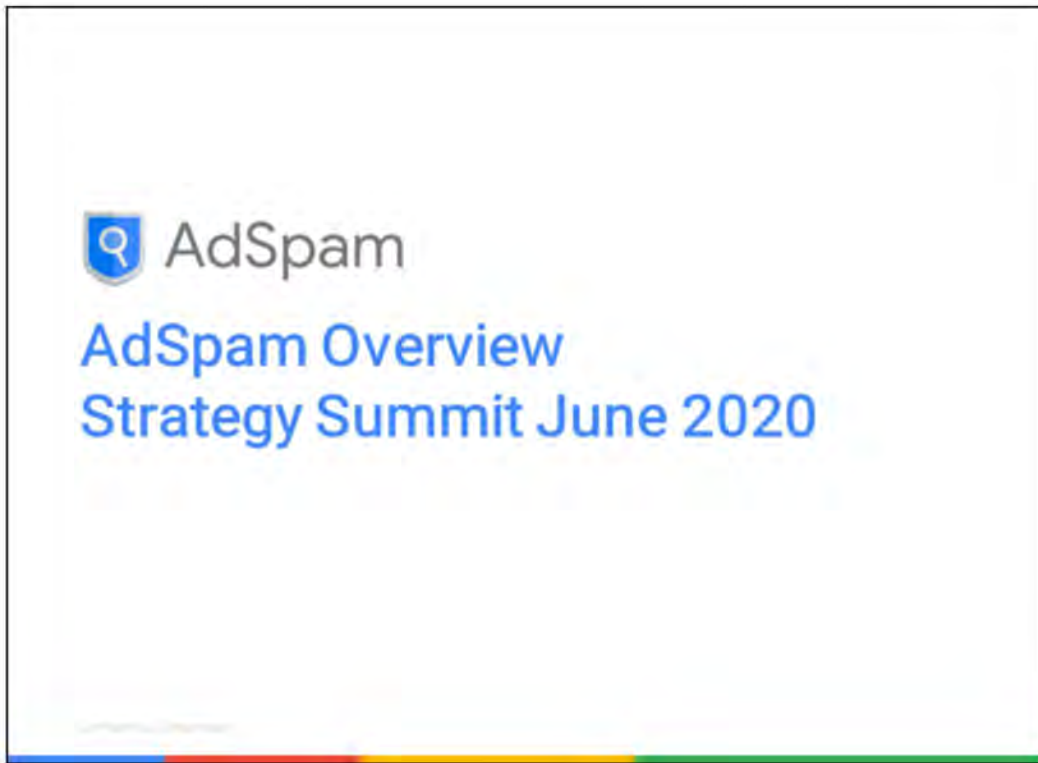
### Logistics Reminder

- Please stick to allocated time. Reminders issued when 5 mins left
- Follow presenter's guidance: asking questions during or after the presentation
- Any required follow ups/additional questions past time - enter them directly to the notes doc
- Note takers - keep AIs and follow up discussion in [go/adspam-strategysummit20-notes](https://go/adspam-strategysummit20-notes)

Google

Misha - 40 sec





Speaker: Subir

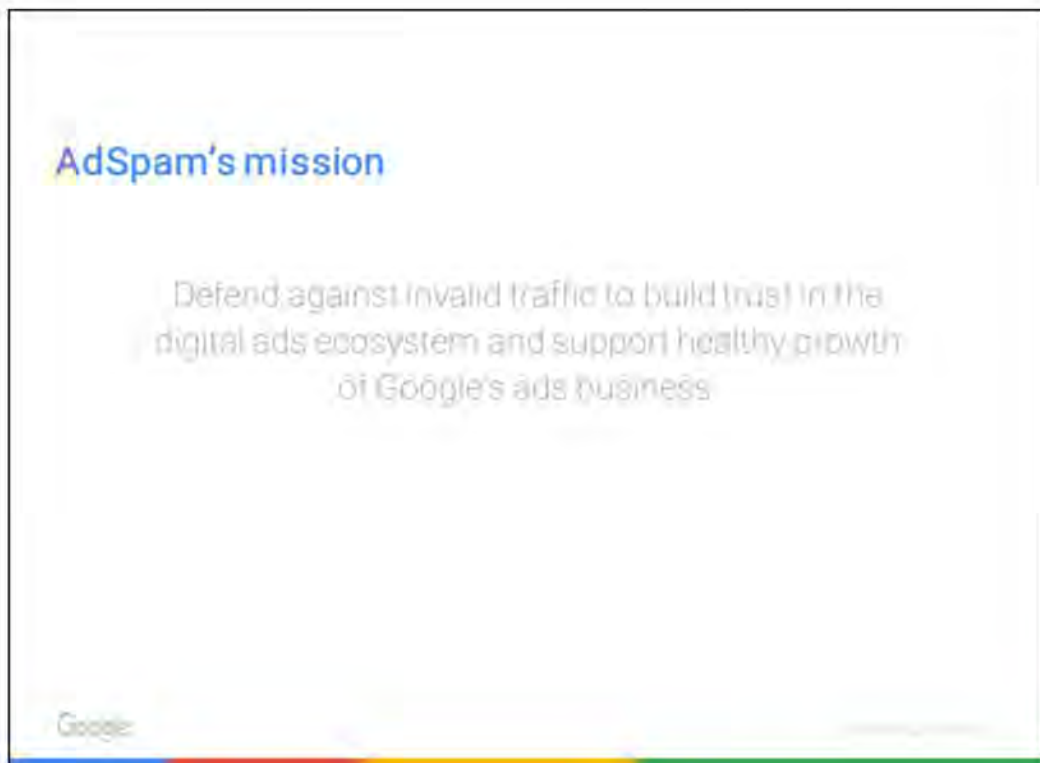





Mission can help answer the big questions...



Alignment on and understanding of key challenges to prepare for 2021 planning (prepare for writing strategy docs for the APaS summit).



The words may be different between how T&S states the mission and how eng states it, but they are both are closely aligned on defending our ads ecosystem against IVT. Puneet explain that T&S' mission closely aligned.



Speaker: Per

Invalid traffic should not be material enough to impact or affect where or how an advertiser allocates budget.

Online advertising ought to get to the same point as some other "mature" areas - for example credit card: As long as you take basic precautions, few credit card holders have material concerns about credit card fraud. Our dream is a future world where advertisers that take basic precautions (e.g., not buy "too cheap to be true" inventory from questionable sources that are not accredited) will not have to worry about invalid traffic and ad fraud in any material way.

Note: As with most other security/fraud areas, we'll "never be completely done". After achieving the "end state" we expect to continue to invest in invalid traffic defenses to maintain that "end state".

Overall guiding principles are available in docs (provide references[1][2])

### Strategic intent (long-term)


Defended and undefended IVT < 1% with less than \$10M confirmed leakage/overfiltration due to client escalations for 1 year

See [go.hrtmagic-intent](https://go.hrtmagic-intent) for an explanation of strategic intent

Digitized by Google

Speaker: Per

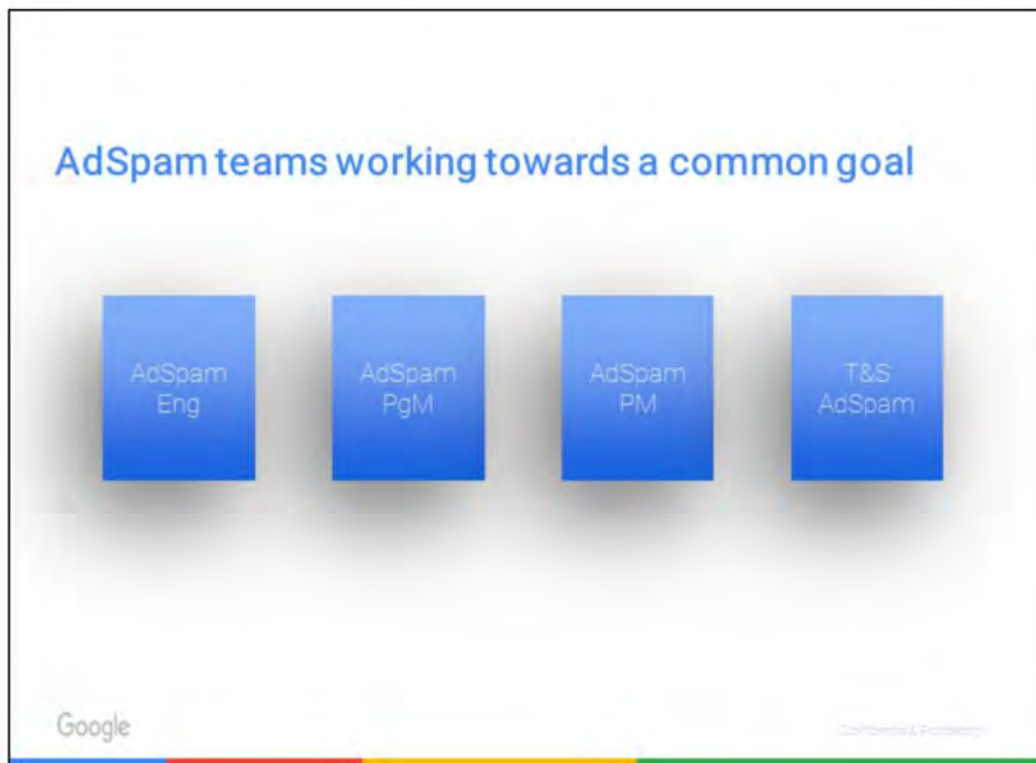
This is a bit of a moonshot. Reminder that SI needs to be specific - know when you reach it. Note that this includes defended (aka caught), undefended (aka uncaught), and unknown (see [go/adspam-framework](#) for definitions) non-policy invalid traffic; however, since we cannot measure/quantify the unknown category we'll only use defended and undefended as the goal.



Speaker: Per

Mention that we'll do the "work breakdown structure" like for SI - define the next level down and each sub-team will also define their respective SIs.



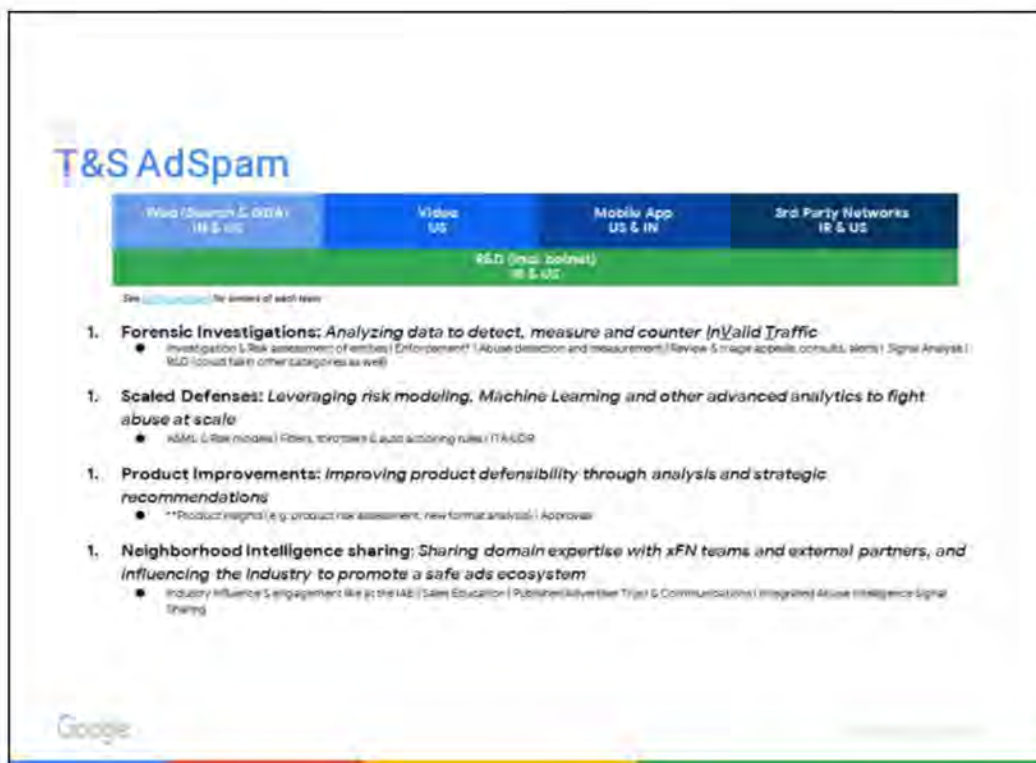


Speaker: Subir





Speaker: Subir



Speaker:Puneet

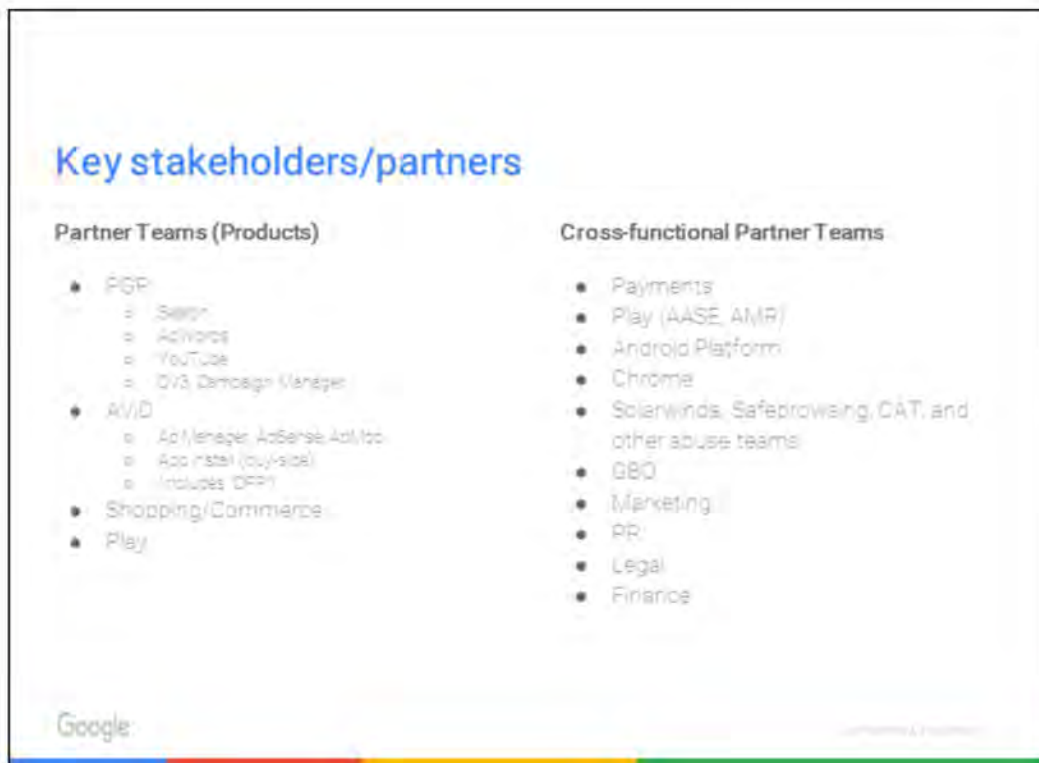
T&S also organizes itself in verticals & horizontals - 4V & 1H. As you can see, there's many similarities how Eng organizes it. Also have project based WGs like Impression spam and a few others.

What T&S does - categories of work...departure from how we used to call Ops and now this new language.

\*Enforcement can include: credits | claw back | termination | zombie | suspension | throttle | warning | reinstate | payment hold | sweeps | blacklist maintenance | contract renewals | workflows / comms | Inbound leads: Appeals | Consults | ML | Governorator | Throttling | research | SafetyNet | Escalations | alerts | Smart Leads).

\*\* Product insights can include: product risk assessment, product roadmap, product strategy, new format analysis, tools, cases, contract review, pub comm

Add'l details on each category available here.




Speaker: Per

## AdSpam's A&C annual goals recap (1/2)

Build trust in Google's advertising metrics by reducing invalid traffic to advertisers, users and publishers

- [Build Defenses] Reduce AdMob SIVT UDR to below 2% (baseline: ~5.4%) while reducing ad fraud on Google Play by \$100M ARR. Projected EOY score: 1.00
- [Network Health] Improve network health by reducing cost-weighted Account Badness Rate (ABR) on AdSense and AdMob by 50% (currently 11% on AdSense and 6.1% % on AdMob). Projected EOY score: 0.70 (75% - network impact on CTR)
- [Impression Defenses] Improve impression defenses to reach parity with click defenses for AdMob, AdSense, AdManager, and YouTube (impression based spend and payments account for about \$x8 ARR). Projected EOY score: 0.70
- [Botnets] Proactively detect >50% of revenue of botnets (incl. "traditional" mobile botnets) that reach "breakout" size in 2020, and reduce average breakout-to-lead latency by 50% to below 90 days. Projected EOY score: 0.70
- [Privacy] Launch issuing TrustTokens and have an approved plan for preserving IVT defensibility without dependencies on 3rd party cookies. Projected EOY score: 0.45



[https://docs.google.com/document/d/1vbl8Vb26gVgeubYqJNWqt41kuTcpODVxWglAf\\_c0ansE/edit#](https://docs.google.com/document/d/1vbl8Vb26gVgeubYqJNWqt41kuTcpODVxWglAf_c0ansE/edit#)

## AdSpam's A&C annual goals recap (2/2)


(cont.)

- [Metrics] Land continuous monitoring dashboards for Undefined IVT Rate (UDR) for core ads product areas: AdSense, AdX, AdMob, DV360, and [STRETCH] Search, YouTube, and GVP. Projected EOY score: 1.00
- [Efficiency] Turn down Blitzzen and Y deprecated systems to return 480K GCUs. Improve initial leads-to-defense time by 50% (to 3 months) and reduce SWE effort on model validation by 50%. Projected EOY score: 0.75
- [Signals] Land Adshield 2.0, the next generation of In-app IVT signal collection libraries, and turn down legacy Adshield 1.0 by EOY. Projected EOY score: 0.70

### Block fraud across PAs using Trust Graph

- [Expand coverage] \$200M impact through AdSec & Adspan/Play collaboration. X% fraud reduction for one ARES client (Cloud, Gala, Gmail, or YouTube). Projected EOY score: 0.75
- [GA] Achieve Trust Graph GA with ASML integration and T&S onboarding. Projected EOY score: 0.60

Google



The slide is titled "T&S AdSpam priorities" in blue text. Above the list are six icons in colored circles, each with a corresponding priority label: a blue circle with a network icon for "Driving cross-product solutions", a yellow circle with a gear icon for "Preventing abuse in design phase", a red circle with a downward arrow icon for "Doubling down on automation for scale", a green circle with a checkmark icon for "Increasing trust and transparency", a black circle with a shield icon for "Cultivating organizational resilience", and a grey circle with a question mark icon for "Excellence in core abuse fighting". Below these icons is a numbered list of seven priorities, each with a brief description and key projects.

### T&S AdSpam priorities

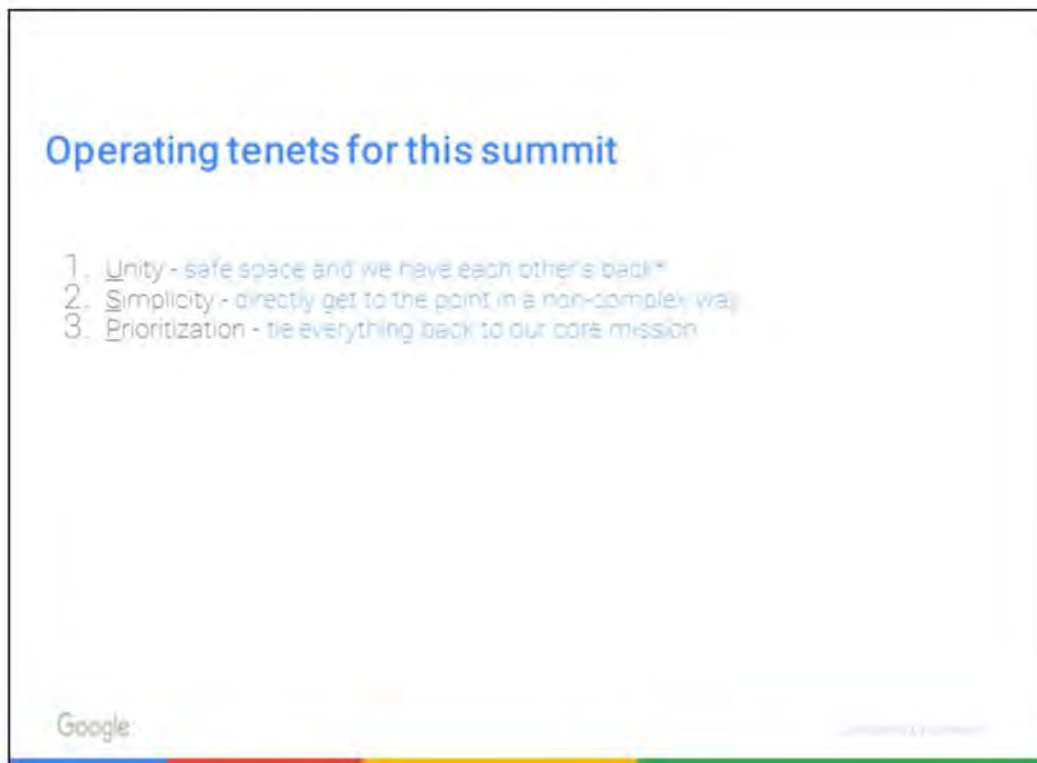
- Driving cross-product solutions**
- Preventing abuse in design phase**
- Doubling down on automation for scale**
- Increasing trust and transparency**
- Cultivating organizational resilience**
- Excellence in core abuse fighting**

- 1. Excellence in core abuse fighting:** Reduce SIVT and escalations with a special focus on high growth verticals of mobile and video | Key Projects: Mobile Ad Fraud ARR reduction, Albus, UDR, impression defenses
- 2. Doubling down on automation for scale & Preventing Abuse in Design Phase:** Find & execute on opportunities to do more with less | Key Projects: ASML, MBL automation, Appeals screening optimization, Organic Labeling
- 3. Increasing trust and transparency:** Improve publisher experience and thereby reduce the legal risk to Google | Key Projects: Greiner, Healthy Apps Ecosystem, Legal Comms
- 4. Driving cross-product solutions:** Partner across AdSpam and Google to reduce unknown and undefended IVT across our networks | Key Projects: R&D, IAI/TG, Privacy, Botnet
- 5. Cultivating org resilience:** Simplify work, career, and work-life harmony | Key Projects: help Sr. ICs build long term careers, make everyone ML adept, 70-20-10 mindset
- 6. Strengthening our Global Operating Model:** Have a rhythm to run T&S AdSpam like a business | Key milestones: summits, newsletters, All Hands, Q&Rs, DMG, etc.
- 7. Evangelizing AdSpam vision:** Drive people towards our vision with our revamped mission | Key Projects: Strong collaboration, speaking "common language", guiding principles/tenets

Google

Speaker Puneet - strategy doc link here

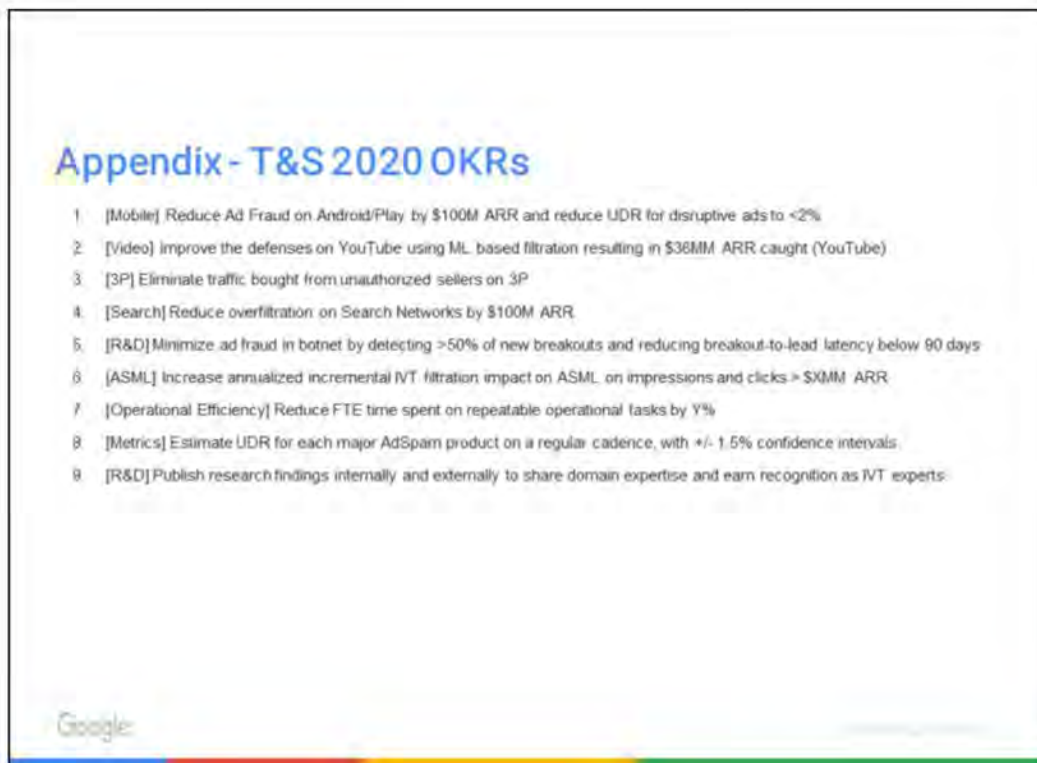
In T&S, we have 7 priorities and aligned with T&S wide 6 pillars.



Speaker: Puneet

\*avoid sensitive ppl issues





Speaker Puneet

Bonus slide if needed - link to OKRs here quickly skim through the slide as most of it should align with common goals.



Proprietary + Confidential

Break ☕ | 10 mins

PST: 8.20 - 8.30 am

GMT: 4.20 - 4.30 pm

IST: 8.50-9.00 pm

Next session: Planning process



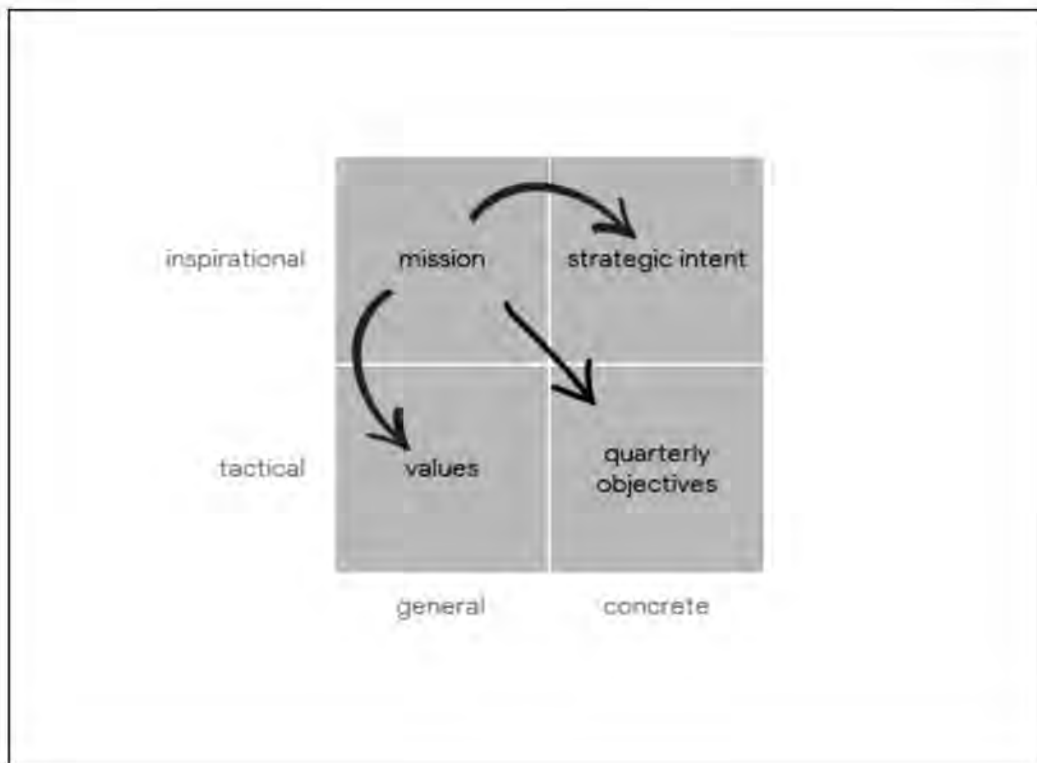
Google

Proprietary + Confidential

# Planning process | 20 mins

PST: 8.30 am  
GMT: 4.30 pm  
IST: 9.00 pm

Google



Mission is critical for driving other elements that we need to run the business

## Strategic Intent

A declaration of an ambition that concretely describes the desired end-state.

*"How will we know when we are done?"*

See [go/strategic-intent](#) for a more detailed explanation of strategic intent.

## Which one do you think is better?

"Develop digital readiness for the transformation of the 21st Century to modernize Britain and create a more digitally-literate country."

Or

"To get everyone in the U.K. online by the end of 2012."

*- Martha Lane Fox, Appointed as UK's Digital Champion*

## **Example: AdSpam Sales Enablement**

100% of our sales (and support) teams are *autonomously* able to communicate our AdSpam efforts with anyone, anywhere, anytime.

*Put differently: zero inbound requests to AdSpam Advocacy XFN team for individual help discussing ad fraud with customers.*

## Planning overview

### Strategic/annual planning (once a year)

- Strategy updates
- Strategic intents (at AdSpam and sub-team levels)
- Annual OKRs (A&C level and sub-team level)

### Quarterly OKR planning (every quarter)

- Team OKRs
- AdSpam A&C level OKRs (small subset of key OKRs)

Strategy updates are directional

Strategic intent is more open ended (can span multiple years)

Annual OKRs are more specific and time-bound?

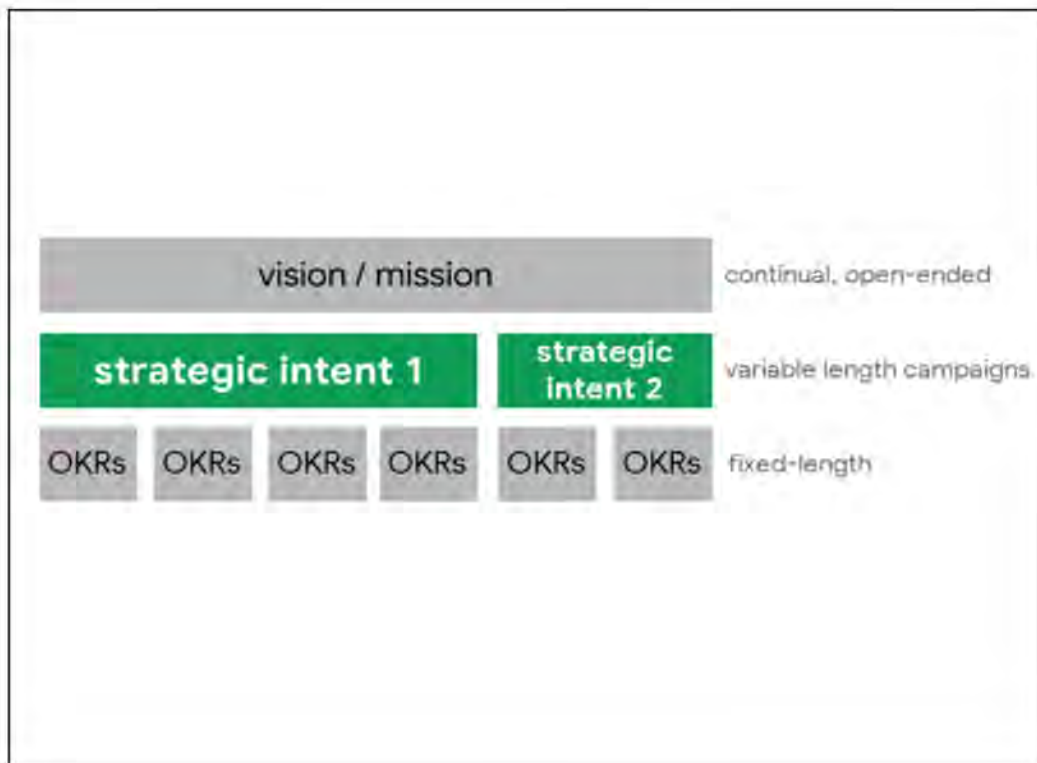
## "Half year OKR planning"

Suggested: Partially follow a "half-year planning" approach

- Half-year (6 months) planning tasks [in Nov/Dec for H1, in June for H2]
  - Gather input from and align with stakeholders
  - Review new trends/strategic changes/etc
  - Stack-rank projects
- Every quarter planning tasks: "official quarterly OKR planning"
  - Estimate efforts
  - Assign resources
  - Above/below the line decisions
  - Enter into EasyOKR etc.

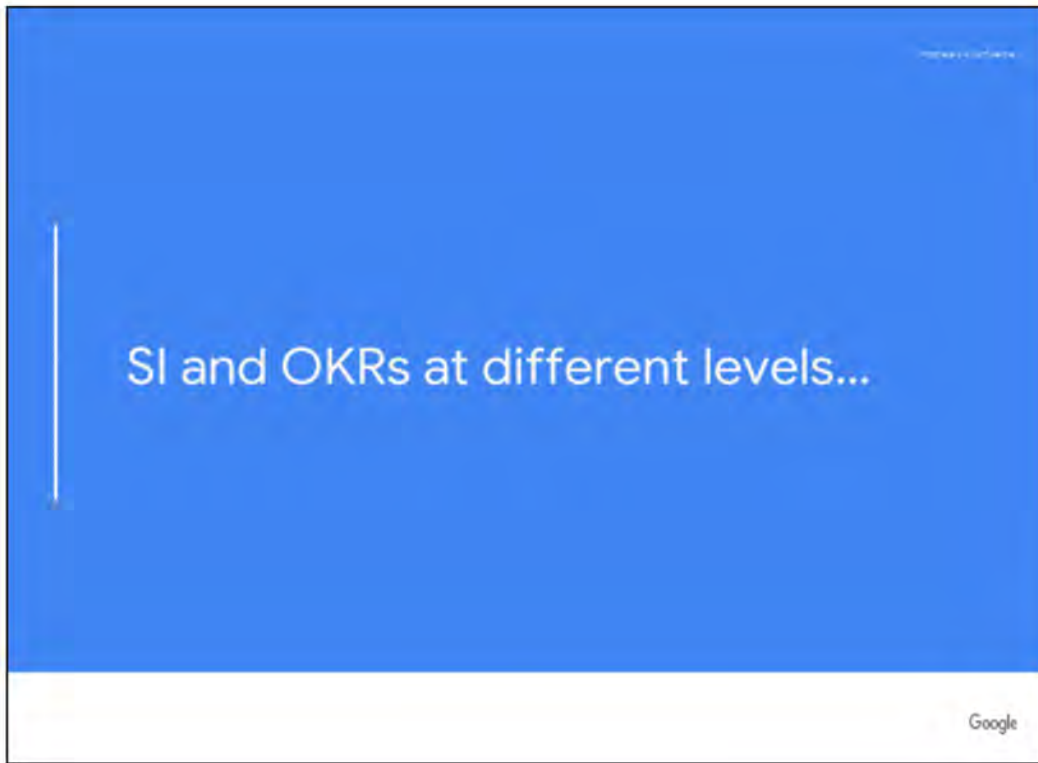
Goal is to reduce the planning overhead in a pragmatic way.

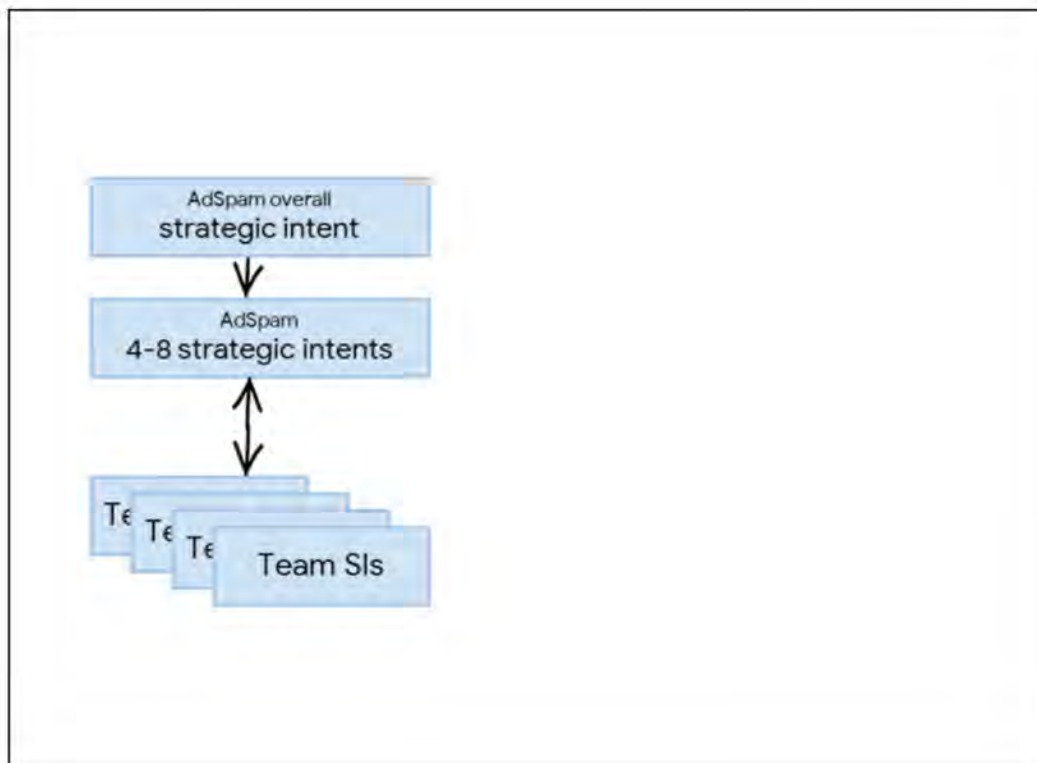





Speaker: Per

While quarterly OKRs are a great way of planning near-term work, the OKRs approach tends to fall short when it comes to strategic planning and management. Why? → Because OKRs are bound to arbitrary durations of time. As a result, teams, organizations, and PAs may find themselves optimizing for the short-term.



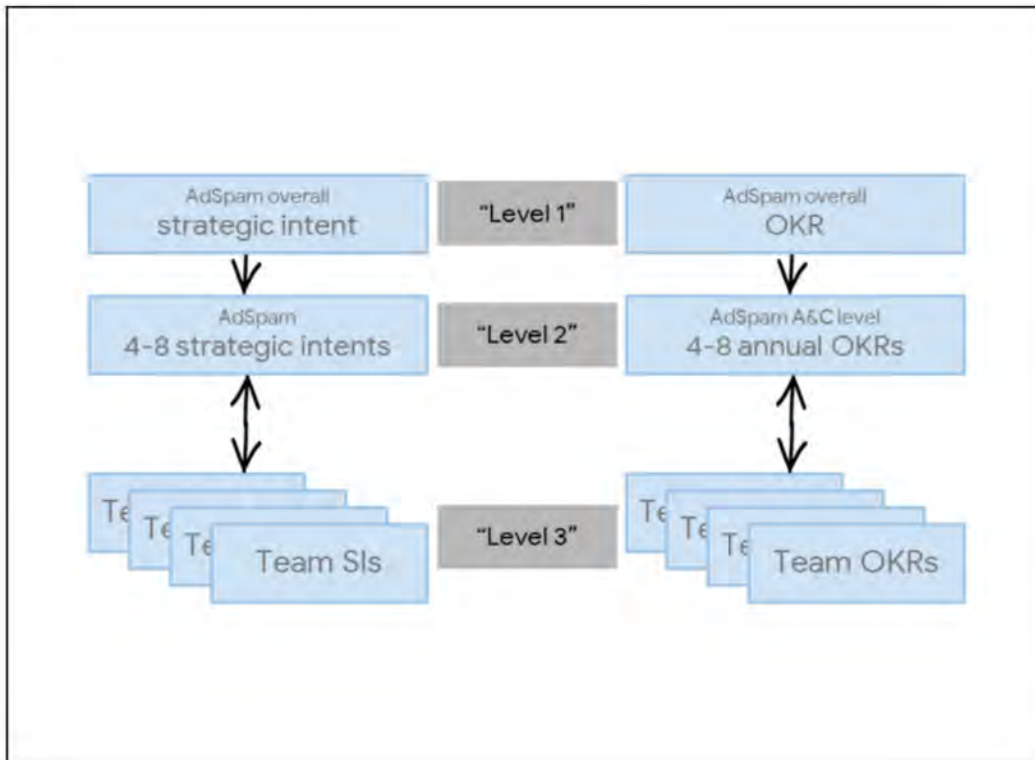


Speaker: Per

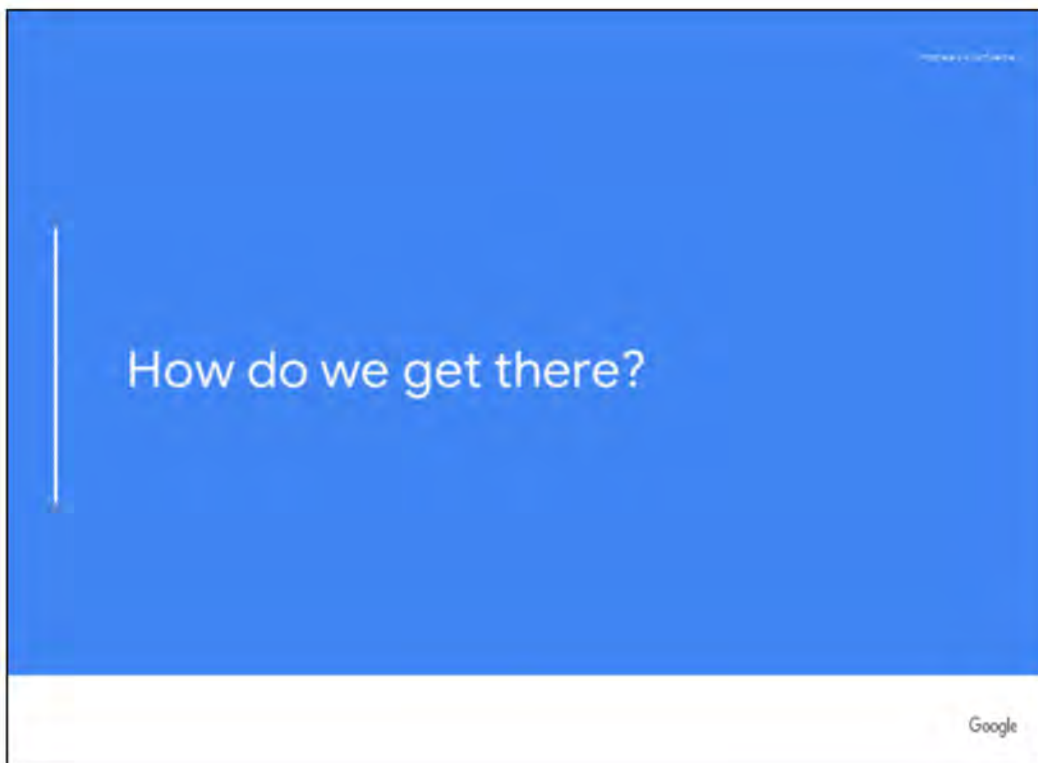


Speaker: Per  
Strategic intents inform/drive OKRs  
OKRs lead teams to realize a strategic intent.

Id	Date	Text
1	06/23/2020 08:45:48	+bjorke@google.com flagging that the relationship should be called out wrt to the Strategic Intents and how they inform/drive the OKRs. Maybe do so verbally or add some text to the arrows. The key point to drive home is that the OKRs lead teams to realize a strategic intent, and not vice versa.



Speaker: Per

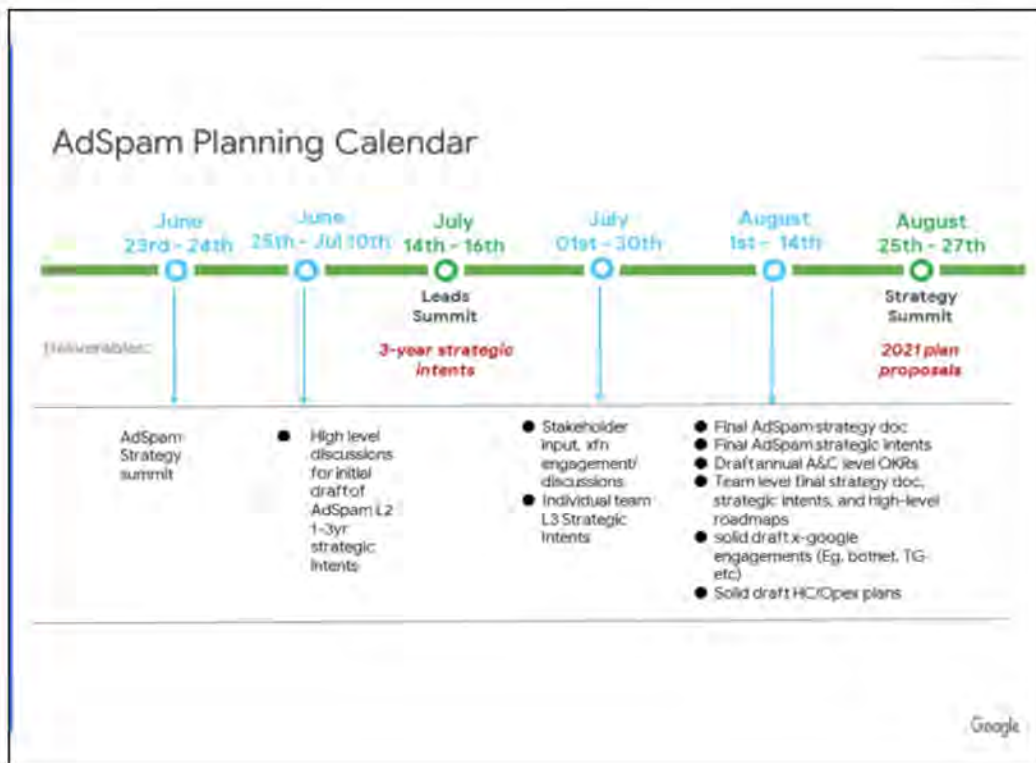




Speaker: Aruna



Id	Date	Text
2	06/23/2020 14:44:17	+bjorke@google.com Small nit: Let's rephrase as high-level strategic intents instead of 3-year strategic intents.
1	06/23/2020 14:44:17	This is APaS level slide as it is. The guidance is for 3 years



Speaker: Aruna

APaS Slides: [https://docs.google.com/presentation/d/1N\\_D989fGKpVq1cuN0UEtFi-EhgzXmyZbxyTLI5kQD4k/edit#slide=id.g5d3751fd3e\\_3\\_53](https://docs.google.com/presentation/d/1N_D989fGKpVq1cuN0UEtFi-EhgzXmyZbxyTLI5kQD4k/edit#slide=id.g5d3751fd3e_3_53)

Draft AdSpam strategy doc

July 14:

Solid draft AdSpam strategic intents (4-7 overall SIs; not sub-team specific)

Team: draft strategy doc, strategic intents (2-5 per team), and high-level roadmap

June/July: Stakeholder input, xfn engagement/discussions, etc.

Aug 25:

Final AdSpam strategy doc

Final AdSpam strategic intents

Draft annual A&C level OKRs

Team level final strategy doc, strategic intents, and high-level roadmaps

solid draft x-google engagements (Eg. botnet, TG etc)

solid draft HC/Opex plans

Nov/Dec:

Finalize annual A&C level OKRs (5-8 OKRs)

Q1 OKR planning

Separately: Q3 OKRs

Leads summit  
3-year strategic intents  
Full doc for pre-read  
1-3 slides to guide presentation  
X-APaS position papers  
1-3 slides to solicit input  
No doc due for 7/14-16

Strategy summit  
2021 plan proposals  
Full doc (draft OKRs, project prioritization, HC / OpEx ask, X-google engagement plans, etc.)  
1-3 slides to guide presentation  
Final X-APaS position papers  
Full doc for pre-read  
1-3 slides to guide presentation  
Note: full X-APaS position papers are due 8/3, but will be shared / iterated up to the strategy summit

Id	Date	Text
3	06/23/2020 09:25:50	+bjorke@google.com same here. Suggest striking 3yr from strategic intents (we don't want to frame as time-bound?)

## X-APaS 2021 Position Papers (AdSpam related subset)

**Future of our content understanding ML infrastructure:** Given the myriad of systems we have across APaS (and PGP/AVE) (around content understanding (images, sites, etc)), what is our Infrastructure strategy to ensure we are not too siloed and reinventing many wheels?

- AdSpam, AI, HUES (shepherd: Rob Malkin)

**Security, Privacy, and On-Device:** We often see privacy and various safety/security efforts at odds with each other (privacy has a natural tendency to want to reduce signal availability, security often depends on more and more signals). We have dabbled in on-device as a possible path to balance these tradeoffs. Given what we have learned in 2020, what is our 3 year vision around transparency, control, and on-device principles?

- AdSpam, HUES, UX (shepherd: Alejandro)

**Team Culture:** Improving our team culture and dynamic is a long-term challenge. What can we do short, medium and long term to improve EDI and empower diversity in the team?

- Mike, Tyese, Reid, Karln, Sublr, TBD? (shepherd: Tyese)

Google

Speaker: Aruna

Id	Date	Text
1	06/23/2020 23:49:35	<p>@aruna@google.com <small>Proprietary + Confidential</small></p> <p>What's Alejandro's Idap?</p> <p>Should we be engaging with the on-device x-apas 2021 position paper?</p> <p>_Reassigned to Aruna Kommu_</p>
2	06/23/2020 23:49:35	<p>aborgia , FYI, Subir mentioned in the summit that he will connect the right group to Alejandro.</p>

## Building for the future

**One-time efforts**

- Platform investments
- Infra/tool migrations (what we know of: Zoomout | New: TBD)
- Others?

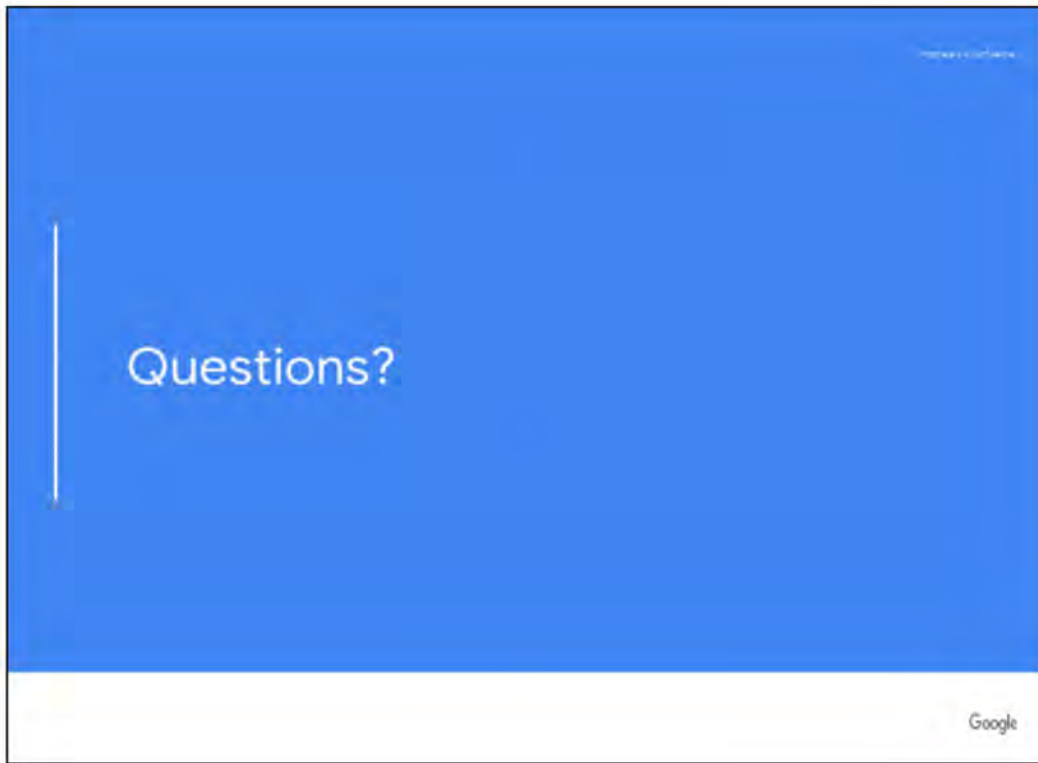
**Recurring efforts**

- Engineering and product excellence efforts
- MRC (2 or so audits per year)
- SOX (every year)
- Resource Planning (yearly & quarterly)

*Reminder: Be empathetic and supportive of asks across sub-teams within AdSpam*

Google

Speaker: Aruna





Proprietary + Confidential

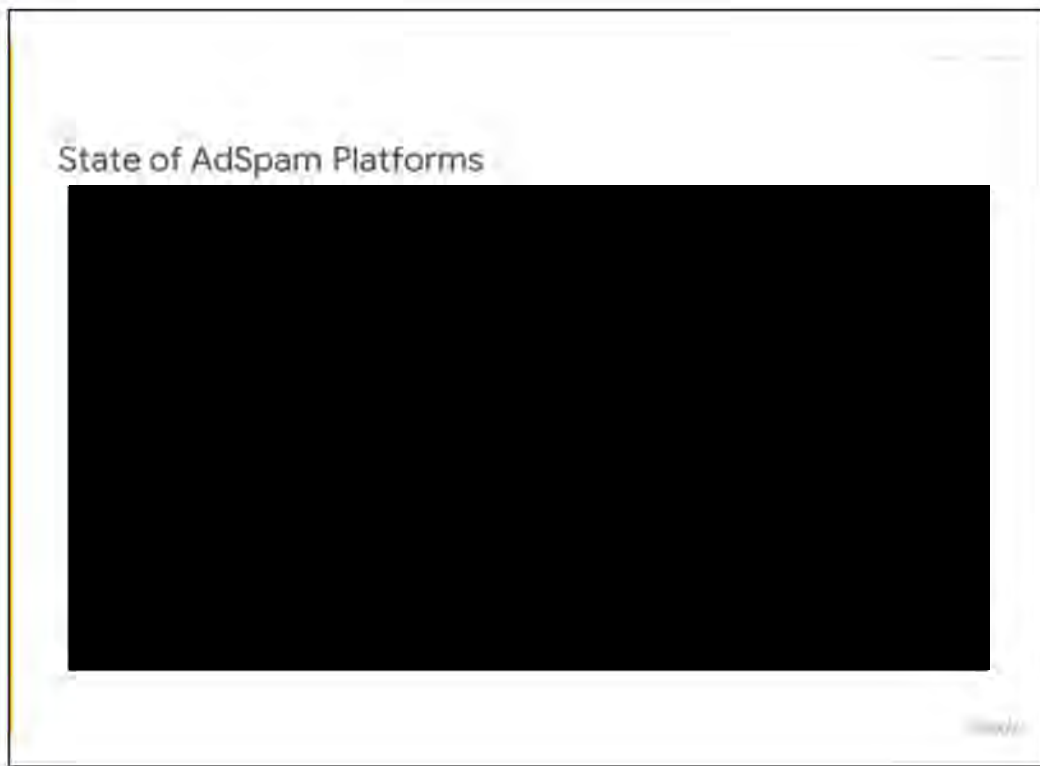
# Future of AdSpam Infra | 60 mins

PST: 8.50 am  
GMT: 4.50 pm  
IST: 9.20 pm

Google

## Outline

- Enforcement and training - prod models, TPU training, prebid and pre-budget, teacher/student training/serving (akademia), launch automation, experiments (rasta integration), lower latency policy
- Turndowns
- Simplifying signal ingestion and data governance.
- Standardizing downstream spam processing - credits, spam/unspam/retractions
- Machine constraints
- Preparing for Privacy, on-device



Famous picture that we use to scare Nooglers in AdSpam. At times this has also helped in convincing senior leadership to do some painful investments. But instead of focusing on the complexity I want to emphasize the strengths.

## Flexibility

- **Training and enforcement** at different feature rich environments - prebid vs post serve.
- **Training options** - aggregation based rules (masonite), weak labellers (LGP/drybell), DNNs (Adbrain)
- **Enforcement options at different timescales** - prebid, pre-budget, post serve, credits/debits
- **Protecting advertiser budgets & users** - pre-bid, pre-budget
- **Enforcement of different IVT** - policy and non-policy.
- **Correction of over/under filtrations** post facto.

Biggest strength is flexibility.

Training & enforcement - joined logs vs just the bare QEM or just the bare CEM

Enforcement options at different timescales - different options have different impact on the business and our advertisers; These also come with tradeoffs around what you can train with.

Protecting advertiser budgets & users against different kinds of spam

### Scaling the platforms - Highlights

- **AsPol (AdSpam Policy)** - The new policy rule enforcement system.
- **Memento** - The new stateless enforcement system.
- **RTAS on all stacks** - CAT2, Click pingback, Call tracking, Viral, Search, DFP, Xbid.
- **ASML enforcement** at click ping back and prebid in addition to post-serve.
- **Teacher/Student training/enforcement** capabilities in ASML.
- **Richer unsupervised techniques** - Representation learning, Active learning, Label propagation, GCN heterogeneous graph processing.
- **Mendel and Rasta Integrations** for streamlined experiments.

Google


AsPol - streaming grouping and aggregation

### Scaling the platforms - Challenges

- **Many more systems** than maintainers - fragmented know-how, monitoring and tooling.
- **High cost** of maintenance - leading to escalations.
- **Cost of tech debt and mandates** keeps increasing.
- **Difficult to reason about** - capabilities, impact, coverage.
- **Inability to invest** in solving other issues - signal ingestion, portability etc.

One of the ways in which we are trying to tackle these challenges is by putting more wood behind fewer arrows and limiting our focus on fewer systems - systems that cover all use cases (policy, non-policy, different enforcement options, different training options) which are built on modern sustainable architectures and technology.

To that end we have an A&c level goal to turn down blitzen which happens to be one of the largest and complicated system in our fleet. Back in 2018 we did a CM that stabilized Blitzen but the cracks are again beginning to show.



Challenges:

- Blitzen non-determinism
- Blitzen lack of past spam capabilities
- No portability
- Incre/decre cannot be measured with metrics

Id	Date	Text
1	06/23/2020 12:08:41	@jmatthews@google.com @vram@google.com sgty? <span>Proprietary - Confidential</span>
1	06/23/2020 12:08:41	Q4 seems very crammed given that its a short quarter. XFP/XFA = 7 pipelines (given that xfp views should get done in Q3).



### Scaling the platforms - Unit Turndowns

Vertical	Unit	Ongoing work	Unlaunch ETA
Search	Dist	Shadow mode	Q3
	Google	Slim review	Q3
	External YJ	pCC and LGP models in soft launch	Q3
	AFS	LGP models in development	Q3
GDA	-	pXY models in development	Q3
Mobile	-	pXY, pTouchDuration	Q3
Video	-	pXY	Q3
	Taco/TVI	Scoping. Not started	Q3/Q4

- Largely making good progress; search & gda -- well defended on asml; mobile & video - pxy;
- unlaunches are followed by shadow mode experiments for a quarter to study if the existing defenses are helping capture a unique fingerprint of spam.
- Shoutout to verticals and TnS - Cross team effort to streamline our defenses and simplify our systems.
- So this is today and how we are walking down the path that we set out in 2018/2019. But whats in store for the future. Beyond stabilizing and simplification we want to invest in areas that can increase AdSpam's impact and improve productivity.

Id	Date	Text
2	06/23/2020 16:47:44	@kendrickb@google.com @jmatthews@google.com does this look fine. I tried to compile info from the status spreadsheet and Kendrick's doc.
1	06/23/2020 16:47:44	For Afs, something like "sprint to develop PCC and LGP models in July" is the current status. Video also has ongoing, required work to launch expanded PCC models.

## Scaling the Platforms - Today & Beyond

### ***Signal management***

- Signal acquisition and plumbing continues to be a large % of verticals workload.
- Standardize production data stores - Woodshed & TG
- Standardize production frameworks for data ingestion - operationally managed workloads - separation of infra failures and data failures.
- Standardize data governance - data dependency tracking, root cause analysis, reduce operational overhead.
- Support easy plumbing and PD analysis of experimental data.

We want you to be able to not worry about how the data is sourced; you should be responsible for validation of sanity of data and the business logic that makes use of this data but everything else should be automated for you.

## Scaling the Platforms - Today & Beyond

### ***Standardizing AdSpam's role in corrections of Ads***

- Does AdSpam's charter include generic corrections in Ads or only IVT corrections?
- Formalize the correction API for Ads - Matcha as a matching API, MaaS for matching, spam/unspam Vs retraction/restatements for correcting:
  - What do advertisers/publishers see?
  - What do downstream teams regard as IVT?
- Move all core stats correction into Conflux.
  - AdSpam only produces spam. Credits/Debits Vs in-month correction all aligned with downstream month close boundaries.

*Google*

This is a broad area that goes beyond Adspam but it is very relevant for our day to day operations. AdSpam is in the business of identify and enforcing IVT - both policy and non-policy. But there have been several incidents in the past where AdSpam annotations are used for correcting experiments that have gone wrong. There are new use cases emerging where Ads looks at AdSpam as a way to correct stats/advertiser payouts.

There are several reasons for this - we have a popular set of APIs and services for picking the relevant events that need to be corrected and we have the systems to produce the annotations that have the intended impact. But are these the correction annotations?

Beyond standardizing and clarifying our role we also want separation of concerns - AdSpam should not be in the business of monetizing IVT/corrections - not an area where our strengths are and we are not properly setup to handle this charter.

## Scaling the Platforms - Today & Beyond

### ***Portability and Upstream defenses***

- Toolkit for portability of defenses at different timescales - prebid, pre-budget, post-serve;
- Automatic monitoring, validation, skew detection and other consistency checks.
- Consistent tooling and evaluation capabilities.
  - Helps protect advertiser budgets.
  - Critical for Real time stats.

### ***Business attribution***

- Build visibility and attribution as a first class use case in AdSpam.
  - Critical to understand impact of our systems
  - Critical for corrections during escalations.

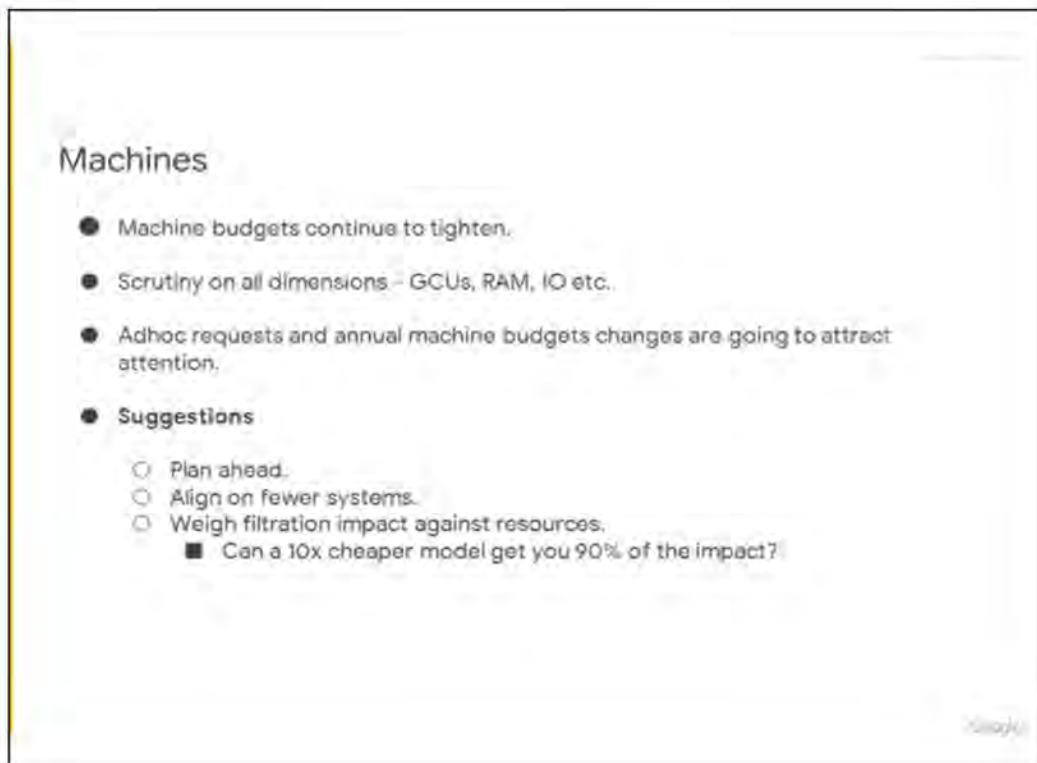
Google

## Scaling the Platforms - Today & Beyond

### **Privacy**


- One of the biggest challenges facing AdSpam today.
- Expected to at least 2-3x the number of variants of filters.
- **Challenges**
  - Scalable what-if analysis and simulation.
  - Scaling our defenses for the long tail.
  - Ease of rollouts/rollbacks.
  - Automatic monitoring and consistency checks.
  - Time bound deliverables.
- Lack of clarity on interplay of on-device learning - is this complimentary to the existing privacy efforts or an additional requirement?

Google

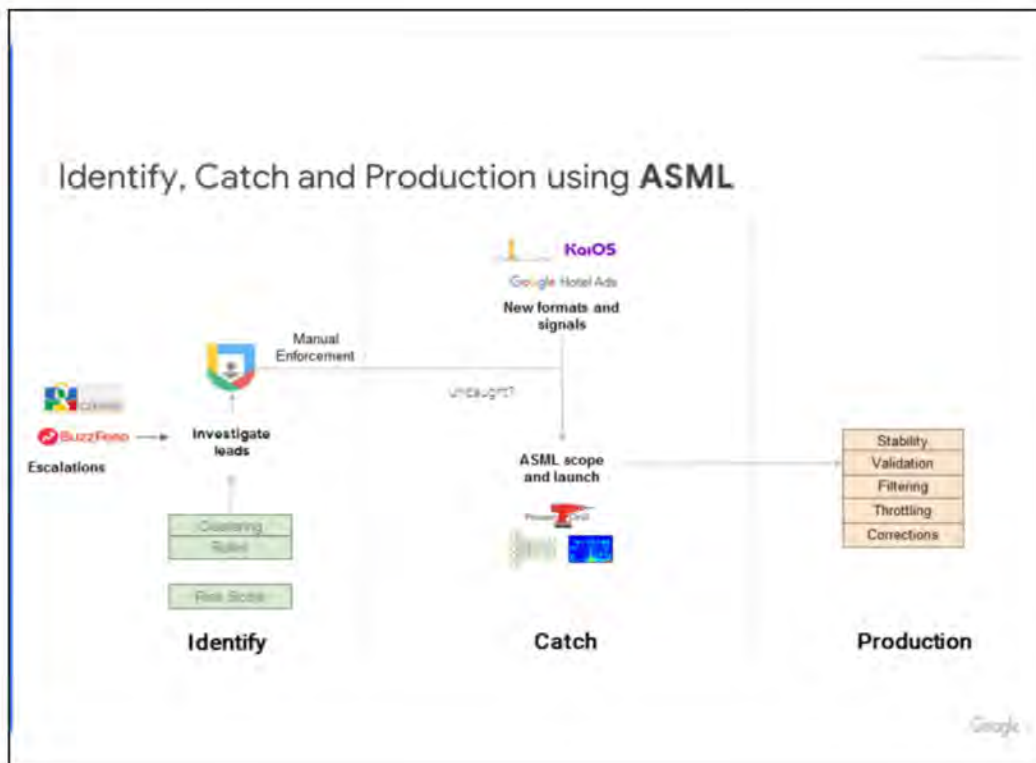


And lastly a word from our sponsors - capital engineering

Ex. display doesn't serve the entire corpus of 300+M ads - only 20M at any time.

Summary: 2018 goals vs 2020 reality 	
Q1 2020 Status	
<b>Coverage</b>	Be the AdSpam solution for 95% non-policy spam
	97%
<b>Faster defenses</b>	
• 90% of spam caught seconds, protecting advertiser budgets	93% pre-budget
• Identify and catch uncaught spam 3X faster	3X faster, \$1.2B incremental
• 3 week launch from spam leak discovery	3 months
<b>Production stability</b>	
• Precision guarantee using online quality monitoring	GoldStone monitoring
• SLA adherence	4 escalations past year
• <2 SWE/Q critical overhead	1 SWE/Q average
<b>Delight the user</b>	
• > 4/5 user rating	4/5
• < 12 hour model iteration	10 hour
<div>Google</div> <div>CONFIDENTIAL - PROPRIETARY</div>	





## Accelerating **Identification** using **Entity Insights**

Automatically annotate each Entity with insights using traffic, content, entity signals

**Anomaly score:** Scores indicating badness with signal-based reasoning

**Related entities:** Entity rings and clusters, entities with same signal fingerprints

**Related badness:** Related to existing or past attack?


**Active learning lift:** How much incremental spam if labeled for ML?

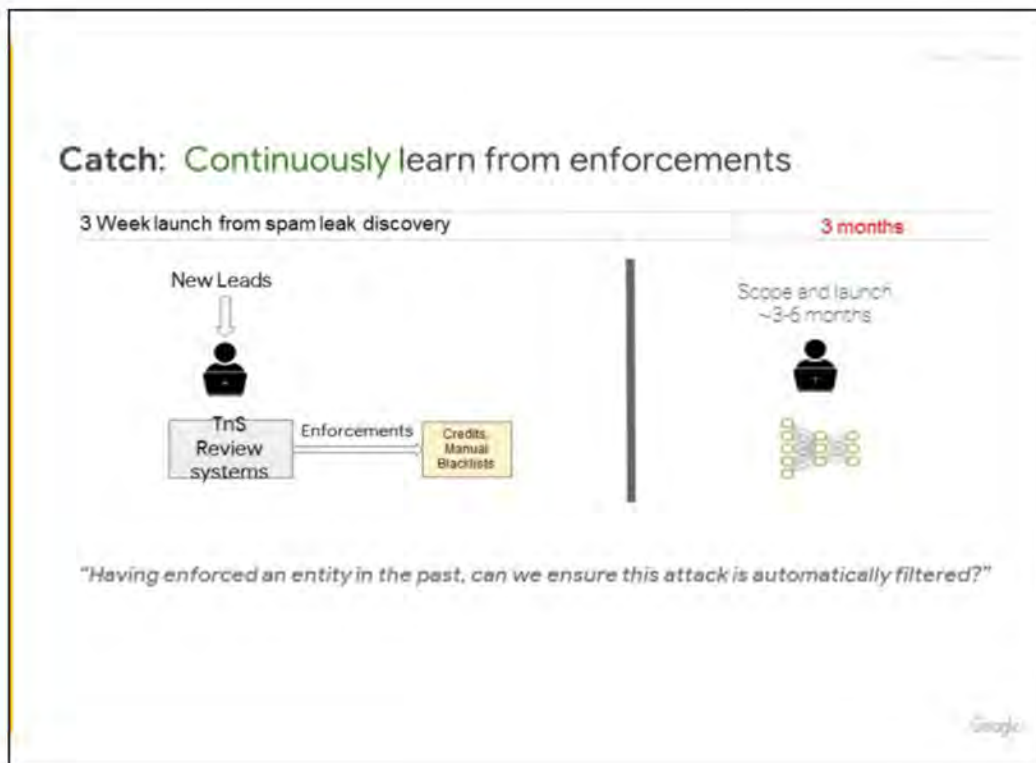
2020 Q2-Q3: Demonstrate feasibility


2020 Q3-Q4: Productionize as part of filter and validation launches

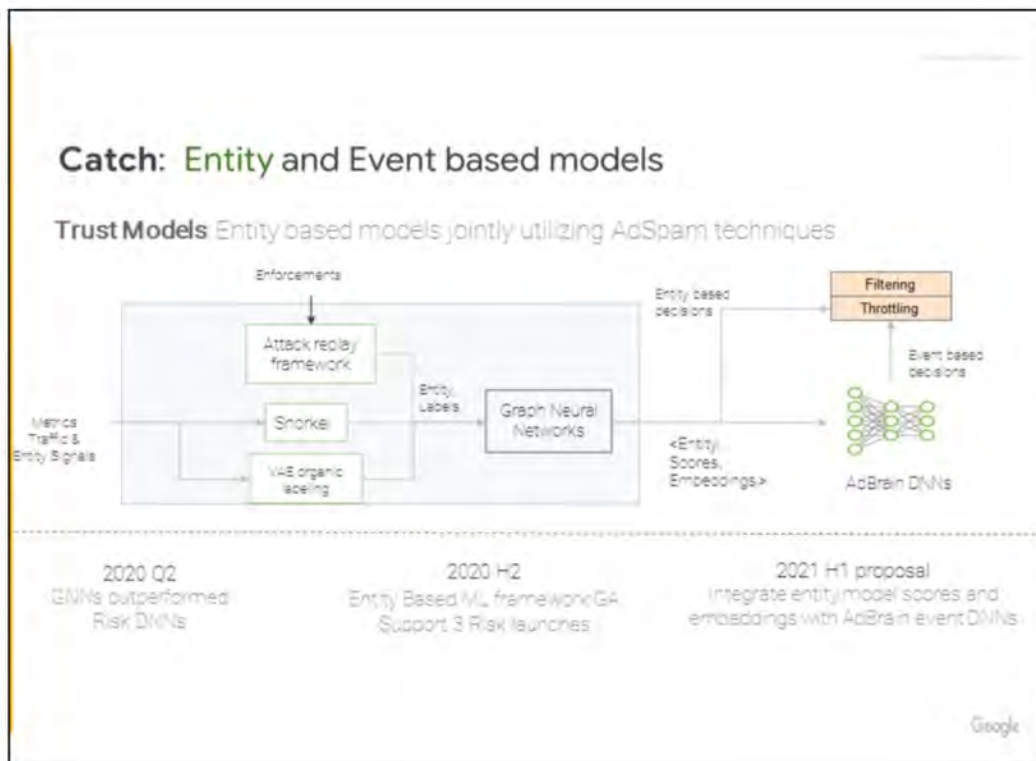
2021 proposal

- H1: GA on all verbiage
- H2: integrate with TrS review tools for **instantaneous** exploration









## Accelerating signal utilization

*"Make it easy to determine if a signal is useful, utilization will follow"*

2019: Manual signal and feature engineering in LBP

2020 H1: Launched VAEs to automatically determine signal buckets, weights, combinations and importance

2020 H2: A/B experimentation framework: GNLvs to utilize both traffic and **entity signals**

(We can accelerate a lot more in 2021!)

### **Signal and rule suggestion framework**

- Rank/order any set of signals based on potential spam importance
- Suggest rules based on signals to generate ground truth

Google

## Transfer learning

*"Create large defenses in areas where signals or labels are lacking"*

Pre-bid				
Impressions				
Clicks				
	GDA	DCLK	Mobile	YouTube

2020 H1  
Transfer learning  
framework launched  
(GDA Q2Q, DCLK pre-bid stopped)

2020 H2  
Support Transfer learning launches  
for impression and pre-bid  
defenses

**2021 proposal**  
Can we strengthen transfer learning to  
defend new formats, surfaces? Privacy-  
safe models?

Google





## Monitoring Gaps

### 2019 Q4:

- GoldStone monitoring limited to ASML
- Handful of metrics used in quality monitoring
- No SFRC; Manual blacklists monitoring
- Quality monitoring seems difficult to explain
- New launches require monitoring configured

### 2020 goal:

*"Any manual or automated filter is automatically quality and quantity monitored with slice-level attribution and 50% actionability"*

Google

## Monitoring Solution

Gaps	2020H1	2020H2
GoldStone monitoring limited to ASL/L	Demonstrate usefulness of GoldStone monitoring on Search and GDA filters	GoldStone to monitor all AdSense filtration
Handful of metrics used in quality monitoring	GoldStone integrated with ASTM. Currently uses 90 metrics	Add all possible metrics (~120). Integrate Botnets metrics
No SPRC, Manual blocklists monitoring		Deploy GoldStone and quantify monitoring for blacklist monitoring
Quantity monitoring spikes difficult to explain	Proof of concept with Graph Mining attributing slices to quantity spikes	Productionize and deploy alongside quantity monitoring
New launches require monitoring configuration changes	Demonstrate Ensemble Metric is adequate to monitor quality of all Search and GDA filters	Automatically monitor <b>all</b> models as soon as they hit live traffic

Google

Id	Date	Text
1	06/22/2020 18:34:23	<div data-bbox="584 183 1058 222"> +ykzhu@google.com please check if this aligns with our plans.  _Assigned to Yong-Kang Zhu_ </div> <div data-bbox="1198 187 1300 198"> Proprietary + Confidential </div>

One last thing... Using spam labels in Pricing models

*"Utilize pSpam to pay the spammers peanuts!"*

Ads pricing model do not use spam information today

- Pro spam and organic similar

Teach pricing model difference between spam and organic (under five experiment)

- Price down spam by 75%
- Repurpose discounts to increase organic spending

New arsenal for AdSpam?

- When in doubt pay spammers less?
  - Today we throttle and consequently lose training labels

The figure consists of two bar charts side-by-side, separated by a blue arrow pointing from left to right. Both charts have a vertical axis labeled '\$' and a horizontal axis with two categories: 'Organic' and 'Spam'. In the left chart, the 'Organic' bar is green and the 'Spam' bar is red; they are of similar height. In the right chart, the 'Organic' bar is green and labeled '+X%' above it, while the 'Spam' bar is red and labeled '-75%' above it, indicating a significant price reduction for spam ads.

Google

Proprietary + Confidential


# Break ☕ | 10 mins

PST: 9.50 am - 10.00 am


GMT: 5.50 pm - 6.00 pm

IST: 10.20 pm - 10.30 pm

Next session: Metrics



Google





## Business Metrics - Introducing UDR Consumer Options

**AdSpam Measurement Options for Vertical Teams:**

- **Undeclared IVT Rate (UDR)** aims to measure AdSpam's hardest question:
  - "How much (AdSense and/or) IVT did we miss?" on a particular product or network
- UDR will take two general forms:
  - **SnapShot** - manually validated, consolidated generally once a year with actionable insights
  - **ContMon** - ongoing measurement of undeclared IVT and lookalikes
- UDR ContMon will have two levels of service:
  - **ContMon Lite** - best effort regular monitoring, requires no additional validation but comes with no warranty of accuracy from Metrics team
  - **ContMon Plus** - best effort ongoing monitoring, requires regular monthly manual validation and comes with Metrics team support for tuning. Manual effort required is smaller sample size than UDR SnapShot, but is optimized to yield highly actionable insights
- Metrics team will look for opportunities to enhance value of UDR ContMon:
  - **ContMon Assist** - Smart tool that suggests manual validation to automate high risk situation mitigated

**Call to Action:**

- We have a RA-level OKR to land UDR measurements. Metrics team is here to support, but can't do it alone
- Verticals need to own and drive development of their risk models - help us hit the goal!



Google

Speaker: Zack

Ads OKR: Land continuous monitoring dashboards for Undeclared IVT Rate (UDR) for core ads product areas: AdSense, AdX, AdMob, DV360, and [STRETCH] Search, YouTube, and GVP.

DTX-951 Page 80 of 181



## Signal Metrics - ASTM (AdSpam Traffic Metrics)

A standardized library of signals and derived metrics:

- Metric definitions are centrally defined implemented, tested, and productionized
- Easily configurable
- Nearly instant
  - No more SQLP → PD slog
- Shared across the team
- Dragged and dropped via Pix Explorer

***"Write once, run everywhere immediately"***

2020 H1	2020 H2	2021 proposal
Integrated with LGR GoldStone TG Assistant 90+ metrics added Complex metric (e.g., CG) infra added	Drag and drop via Pix Explorer Integrate with TG Storage Onboard Botnets pipelines Autotuning infra (non-CG)	Retire onboard all classic pipelines Integrate with AdBrain Integrate with TrS review tools

Google

Speaker: Souvik

## Signal Metrics - Realtime Auto-slicing and Ensemble Score

**Goal:** Reduce manual validation effort and make validation consistent across launches.


**Ensemble score:** Aggregate all (weak) metrics into a strong ensemble score, emulating what a human will do to validate spam (Inspired by [Snuba](#)). Search, GDA POC completed. Mobile next.

### Current filter validation flow

1. Choose a certain entity (on GDA, in most cases it is WP x SubWP x Domain)
2. Manually verify top 50-100 entities such that the top 50% of the incremental traffic of the model is verified.
3. Randomly choose and verify another 50 entities from the long tail, to ensure that the model has no inherent bias.

### 2020H2 filter validation flow

1. Utilize [Certification](#) slices, rather than top entities, because it finds larger, yet coherent, slices. ~40% less slices to validate
2. Ensemble score labels each slice as either SPAM, HAM or Borderline
3. [Importance Sampling](#) chooses top entities from the three above classes

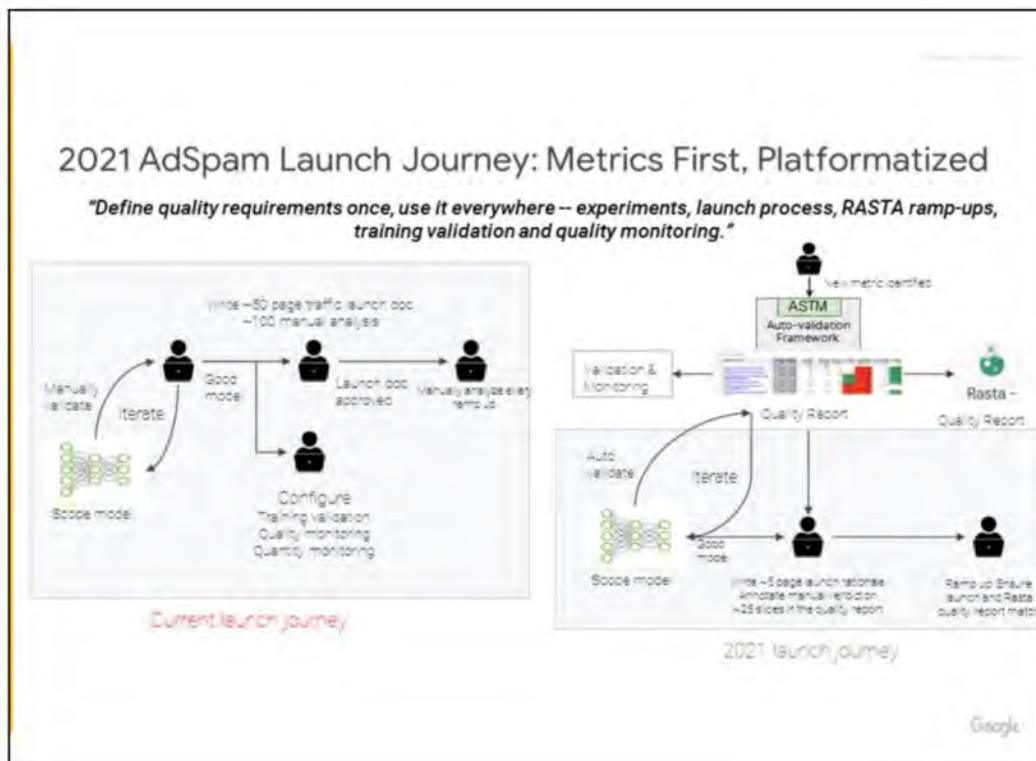


1. Manual analysis mostly on Borderline slices.

Expect 60% reduction in manual reviews to maintain the same quality checks needed for launches. Each slice annotated by all metrics and ensemble score, improving consistency.

*Google*

Speaker: Souvik



Speaker: Souvik

Proprietary + Confidential

Break ☕ | 10 mins

PST: 10.30 - 10.40am

GMT: 6.30 - 6.40 pm

IST: 11.00 - 11.10 pm

Next session: Privacy



Google

Proprietary - Confidential

Break ☕ | 15 mins

PST: 10.30 - 10.45am


GMT: 6.30 - 6.45 pm

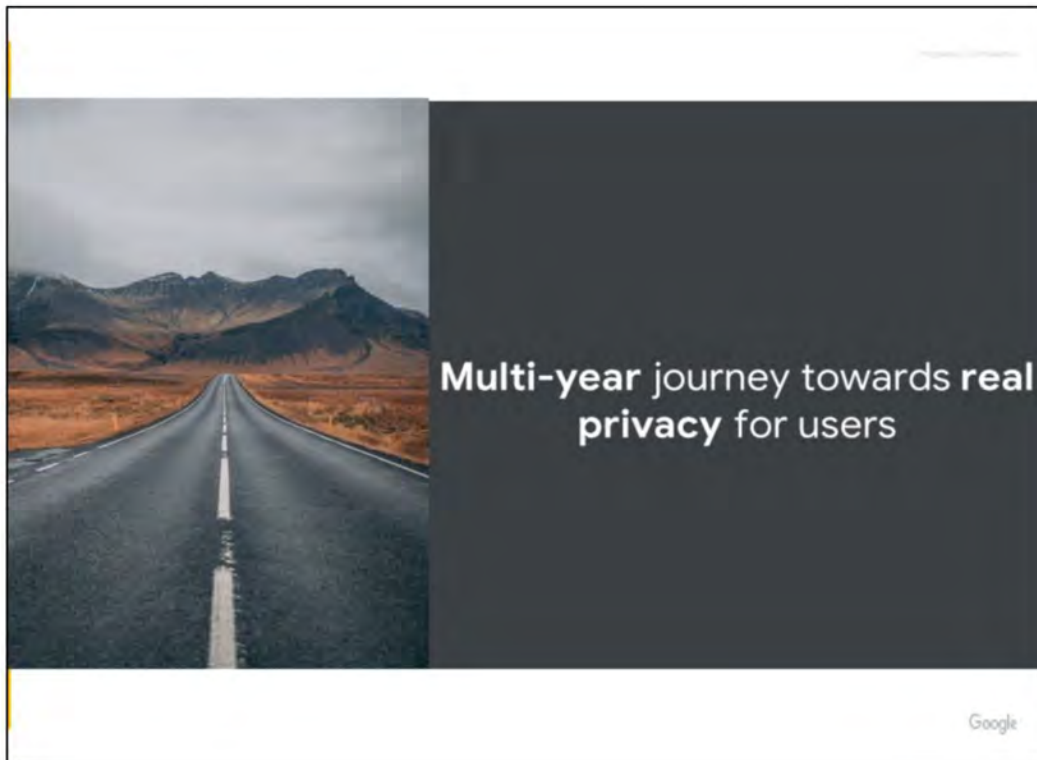
IST: 11.00 - 11.15 pm

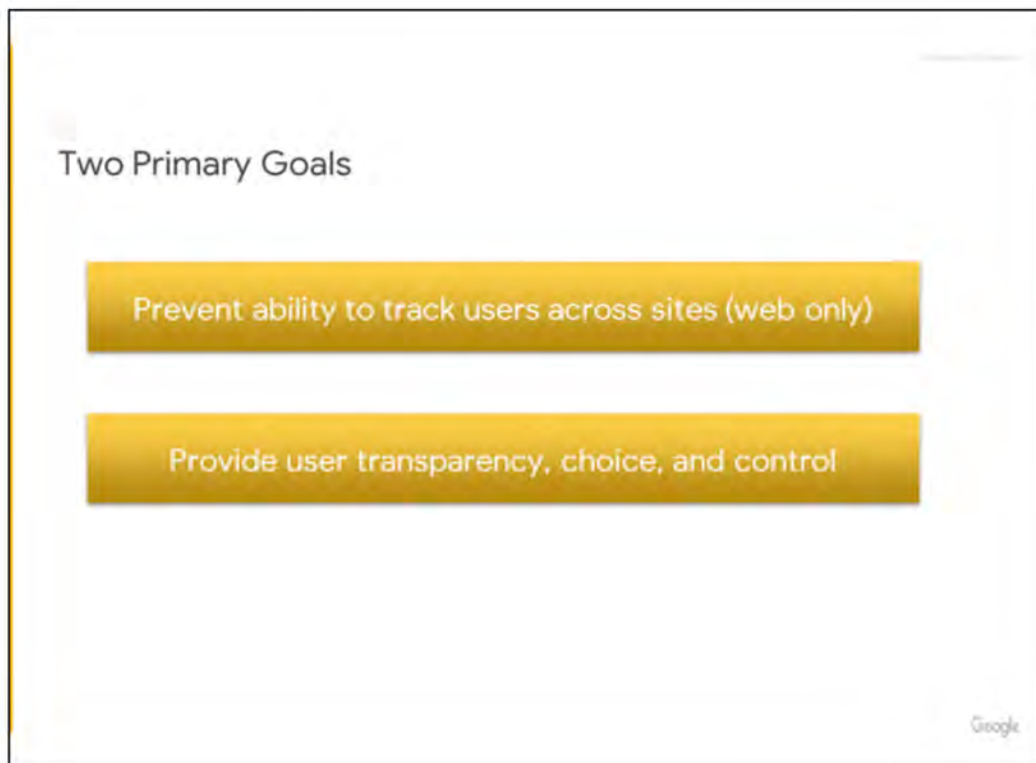
Next session: Privacy



Google









## Major Privacy Initiatives

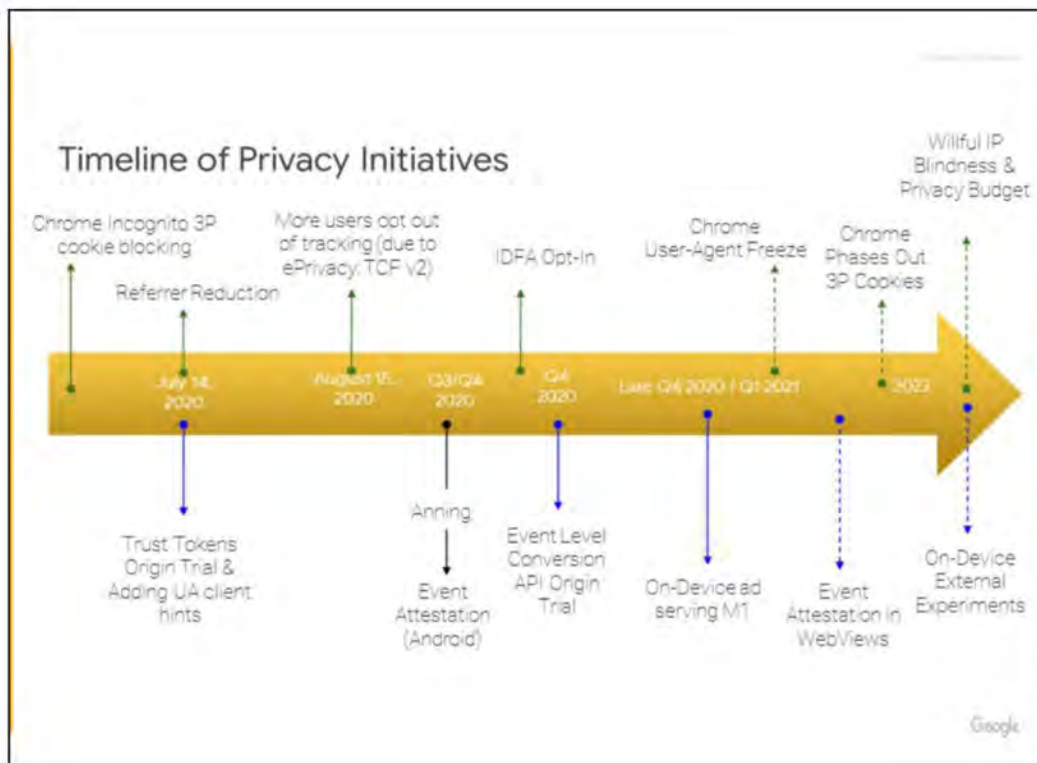
- **Web**
  - **Chrome's Potassium:** Prevent ability to track users across sites (Privacy Sandbox)
  - **Ads' Newton:** User transparency, choice, and control (opt in/out, who is collecting, delete)
- **Mobile apps**
  - **Anning:** User transparency, choice, and control (opt in/out, who is collecting, delete)
  - **Event Attestation (aka "Secure ADID"):** Core technology to make Anning work
  - **On Device Ads:** "No data should leave your device"
- **Consent regulations leading to cookieless ads (ePrivacy etc.) (web + mobile)**
  - No ads cookie/ADID etc. without user consent (currently in EU)

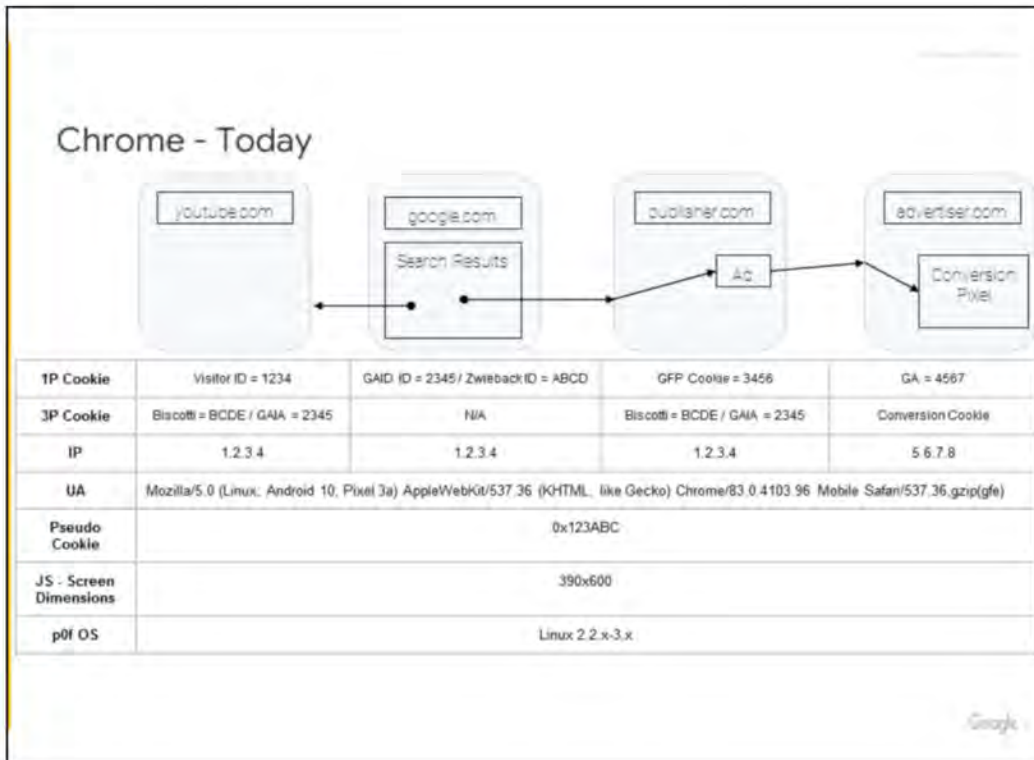
Google

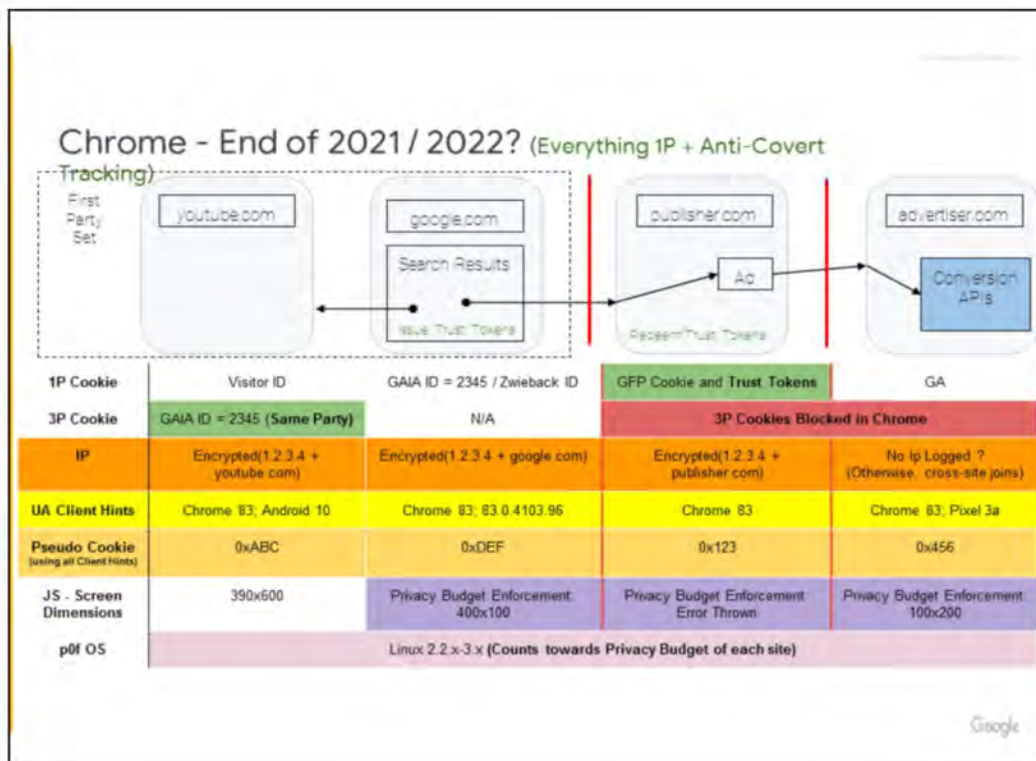
## What Does This Mean For AdSpam?

- **Privacy for web:** Chrome's "Potassium" + Ads' "Newton" projects
  - How can we minimize the loss of defenses with no 3P Cookies + Anti-Covert Tracking?
- **Privacy for mobile apps:**
  - Anning + **Event Attestation**: WebViews vs In-App Browsers
    - Different defenses for Display/Video ad traffic within WebViews in Anning?
  - **On Device Ads**
    - No more logging of individual events (e.g. no more QEM, CEM)
- **Web + mobile: Consent regulations leading to cookieless ads (ePrivacy etc.)**
  - How do we handle loss of 1P/3P cookies/AAID with ePrivacy?

Google







KEY TAKEAWAY: Can't track users cross-site (only within same "company")

## Research we need to explore?

- How can we achieve cross-site based filters (e.g. co-click) with help from Chrome?
  - Chrome's proposed [Aggregated Reporting API](#)
  - FLoC IDs?
  - Bloom filters in Chrome with Differential Privacy?
  - Crypto (Secure Multi-party Computation) + Cloud ML?
- Anti-abuse teams want IPs, but Chrome wants us to move away from IPs
  - We're always asked if we can rely less on fingerprinting and more on Trust Tokens
- Can we build defenses only on 1P information with privacy budget enforcement?
- How well will CG (Q3) and Conversion Spam work with Differential Privacy?
- Optimal 1P Trust Token Issuance logic and developing 3rd party issuers (e.g. whiteops tokens)
- How much entropy does AdSpam need w.r.t a privacy budget?
- Effectiveness of Albus and third party IVT detection within Chrome Privacy Sandbox?

Google

## ePrivacy & Cookieless ads (August 2020)

- TCF (Transparency & Consent): industry framework for communicating user *consent*
  - TCFv2 cut-off date August 15, 2020 and more user controls.
- **ePrivacy Directive**: EU regulation that requires *consent* to use cookie or similar technology on user's device for non-essential purposes
- Serving cookieless ads:

Narnia 3 (O&O)	Display	Conversions
<ul style="list-style-type: none"><li>• Zweback/Visitor ID is a "mixed" use cookie (essential and non-essential)</li><li>• Proposal: New AdSpam only cookie or encrypted Zweback/Visitor (Q3 2020)</li></ul>	<ul style="list-style-type: none"><li>• Ads related cookies are not essential.</li><li>• Serve cookieless ads on reservation and app mediation traffic.</li></ul>	<ul style="list-style-type: none"><li>• Conversion tracking needs consent.</li><li>• <a href="#">Project Macaron Code Yellow</a></li></ul>

- AdSpam implications:
  - Short term: no cookie based defenses on cookieless traffic
  - Longer term:
    - Modifications to Event Attestation / Trust Token to remove public key
    - On-Device & In-Memory Frequency filtering

Google

## EventAttestation and WebViews (Q4 2020 -> 2021)

### Issues with Hybrid apps

- WebViews have their own cookie jars and separate from AAID/IDFA
  - AdSpam developed Biski to address this
- User Privacy and Transparency controls?
- Ambiguous if dealing with an 'App' or 'Browser'

### With Anning - App developers to choose if they are more 'App' or a 'Browser':

- If 'App', then WebView can receive EventAttestation
  - Secure AAID + Basic DroidGuard attestation + Confirmed WebView!
- If 'Browser', then (long term) WebView removes 3P cookies and should use Chrome's Privacy Sandbox APIs

### Open questions:

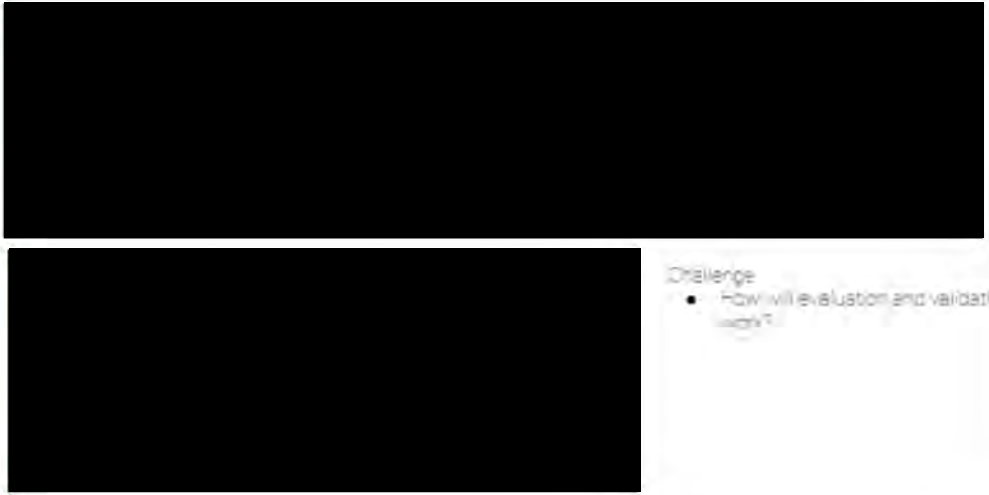
- GDA fragmentation? GDA web and GDA app?
- Is AGSA or Facebook app more 'app' or 'browser'?



## On-Device Ads (Project [Leibniz](#)) and AdSpam

- M1 milestone - AdMob/DRX app inventory, GDA interstitial + DV3 video demand
- Opportunities for *new personalization signals*, but many challenges:
  - No more event level logging (QueryEventMessage, ClickEventMessage)
    - How will past spam, recoveries, credits work in this new paradigm?
    - Pre-bid and post-serve filtration?
  - No more GAIA/AAID/Biscotti ⇒ Cookie Frequency defenses move "on device"
  - Train pModels and combiners on device
  - Graph building (LGP or co-click) on device?
  - New threats:
    - Compromised devices stealing the models or tampering with the weights
    - Access user/publisher/advertiser data
    - Reverse engineer on-device auction

On-device Spam Fighting with ML: Options




Challenge

- How will evaluation and validation work?

Google

## On-device Spam Fighting with ML: Options

*"What if we have access to no events server-side, only aggregates?"*



- Per-event prediction and enforcement on-device
- **Several challenges**
  - AdSpam Specific
    - How to build meaningful base model using only aggregated features?
    - How will model monitoring work under differential privacy?
    - Several techniques to "fine-tune". Which technique? How to validate without access to data?
  - General
    - How to prevent stealing and reverse engineering of model and enforcement mechanisms
    - Unclear serving feasibility at AdSpam scale
    - On-device model size, complexity constrained by memory and battery budget.

Google

## Fragmentation of Defenses

- Chrome and other Browsers
  - Will other Browsers adopt Chrome's Privacy Sandbox?
- Chrome Android potentially having more signals through GMScore
  - e.g. Android Trust Tokens; perhaps on-device integration?
- Browsers vs WebViews
  - Display/Video ads in Browsers (no 3P cookie) vs in WebViews (Injected Event Attestation)
- Conversions
  - non-Chrome Conversions vs Chrome Conversions using Differential Privacy
- Browsers (limited entities) vs Apps (AAID/IDFA and IPs)
- ASML and Android GMScore on-device

Google

Necessary "remedies", but leads to more complexity?

## Call To Action

Privacy changes affecting AdSpam are starting **now!**


Be mindful of privacy changes that's coming within ~2 years vs longer term efforts (e.g. on device).


Do we continue to launch web-based defenses using entities/signals that we know are going away within 2 years (e.g., 3rd party cookies, detailed UA and IP, other high-entropy signals)?


**Request teams start investing time in new web-based defenses using alternative signals and approaches.**

*Note: Privacy Team has limited resources with ~3 SWEs and 1 T&S*

*Google*









## Agenda

### Day 2: Deep dives into key challenges

Keynote by Salkat
High level strategy overview
break
Trust Graph powered analysis
break
Challenges with supporting new/emerging markets/ segments/ products
Improving precision and monitoring
break
Impression Spam
Apps & Play
Closing Remarks
Fun event

See detailed agenda at <https://adspecs-strategysummit20-agenda>

Google

Proprietary + Confidential

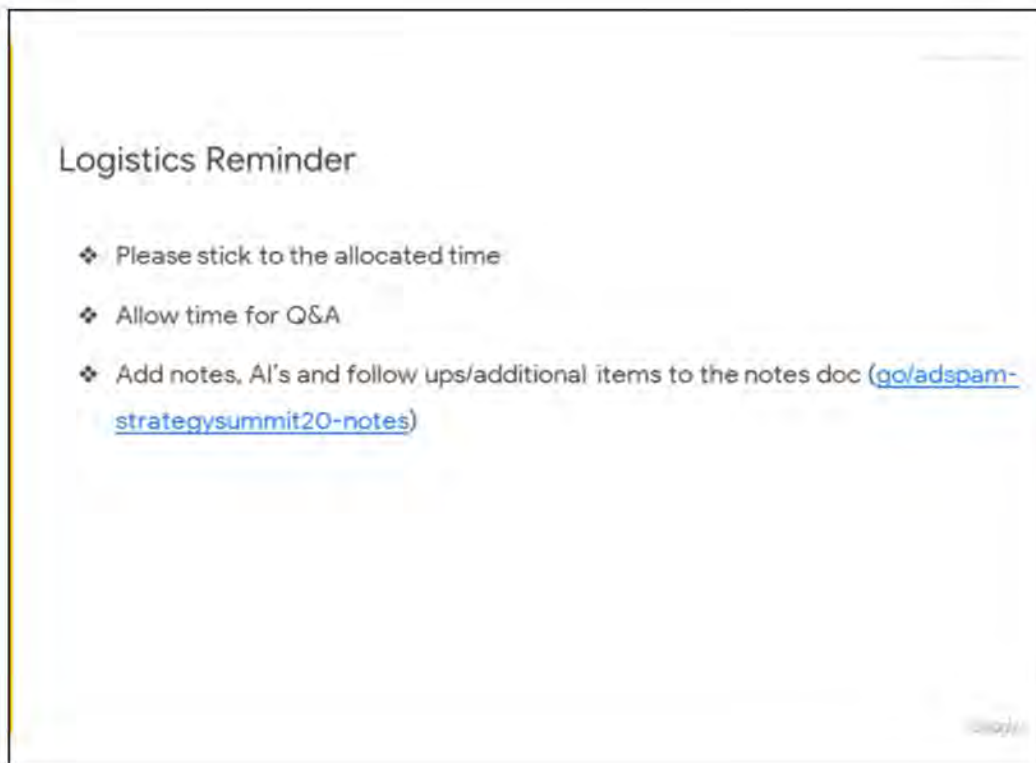
# Keynote by Saikat

PST: 7:30am - 7:45am

GMT: 3:30pm - 3:45pm

IST: 8:00pm - 8:15pm

Google



Misha - 40 sec

Proprietary + Confidential

# High level strategy overview | 15 mins

PST: 7:45am - 8:00am

GMT: 3:45pm - 4:00pm

IST: 8:15pm - 8:30pm

Google

Id	Date	Text
1	06/24/2020 04:49:32	+andresf@google.com +subir@google.com +aruna@google.com Added the strategy slides -> will check early Wed morning if you have suggestions for improvements.

## Session goals

- Update on progress towards updating the overall AdSense strategy
- Next steps (call to action)

Google



New platform/infra allows us to take it to the next level - what innovations can we build on top of our infra?

We are developing new and promising approaches - clustering, risk modeling, organic labeling, etc.

We are getting much better at leveraging data across Google

Industry collaboration, standards, and hardening the ecosystem is gaining momentum

I think we are ready for take-off to new levels of IVT defenses.


## Structure of the strategy doc

- Introduction: Summary, background, definitions
- Goals
  - Brief description
  - Key challenges
  - Gaps/risks and opportunities
- Strategic intents
  - Core projects & roadmap
- Major dependencies and risks

Google

Confidential & Proprietary





**1) Close gaps in IVT defenses and address largest IVT risks**

Minimize IVT risk by addressing areas with weak defenses across Google's ads products, and address the largest IVT risks.

**2) Ensuring long-term sustainable IVT capabilities: Build for the future**

Ensure that AdSpam is able to effectively protect Google's ads business in the future: win the arms race with bad actors, adapt to ongoing shifts in privacy approaches and other ecosystem changes, and evolve in tandem with Google's ads products.

**3) Support safe growth of new ads business**

Support the sustainable growth of new business opportunities like new ad products, new formats, new environments, new deals, etc. across all of Google's ad product portfolio.

**4) Protect Google by increasing IVT defenses and trust across the industry**

In order to effectively fight and prevent IVT and ad fraud within Google, we need to increase IVT defenses and trust across the industry at large to minimize overall systemic risk.

## 1) Close gaps in IVT defenses and address largest IVT risks

**Description:** Minimize IVT risk by addressing areas with weak defenses across Google's ads products, and address the largest IVT risks.

**Challenges:**

- Large number of surfaces to defend (products, formats, devices, event types, logs, etc.)
- Adversarial nature means we can't launch in a "set it and forget it" mode
- Sophisticated fraud is attractive for some companies and cyber criminals
  - They are highly skilled and can mount sophisticated and scalable IVT operations (attacks)

Google

## 2) Ensure long-term sustainable IVT capabilities: Build for the future

**Description:** Ensure that AdSpam is able to effectively protect Google's ads business in the future: win the arms race with bad actors, adapt to ongoing shifts in privacy approaches and other ecosystem changes, and evolve in tandem with Google's ads products.

**Challenges:** Developing a deeper understanding, more comprehensive knowledge, and better methods, tools, and techniques to

- scale with Google's evolving ads products
- adapt to dynamic and increasingly sophisticated bad actors
- maintain defenses in the new privacy-first online environment.

Google

### 3) Support safe growth of new ads business

**Description:** Support the sustainable growth of new business opportunities like new ad products, new formats, new environments, new deals, etc. across all of Google's ad product portfolio.

**Challenges:**

- Providing IVT defenses for new products and deals that may not "fit the mold"
  - new interaction models
  - not using "standard" ads backends (e.g., logs)
- Supporting custom deals while maintaining economy of scale and avoiding one-off solutions
- Achieving defensibility in environments that lack signal collection

Google

#### 4) Protect Google by increasing IVT defenses and trust across the industry

**Description:** In order to effectively fight and prevent IVT and ad fraud within Google, we need to increase IVT defenses and trust across the industry at large to minimize overall systemic risk.

**Challenges:**

- Sharing information without
  - leaking privacy sensitive data
  - giving the secrets to "bad guys"
- Hard to motivate an industry mostly focused on short-term revenue to invest for the long-term sustainability
- Lack of ground truth/effective advertising ROI metrics makes it hard to justify investments
- Monetary incentives not always aligned with sustained and effective IVT defenses

Google

## Topics to consider for L2 strategic intents

Some ideas/food for thought:

- Not relying on 3rd party cookies for AdSpam defenses
- "Connect the dots" leveraging data from across Google (Trust Graph etc)
- Impression defenses on par with click defenses
- Know your users and partners (entities, pubs, supply partners, sites, apps, etc.)
- SIVT level defenses for all DCLK products
- Productivity/time from attack to launch
- Use organic labeling across all areas
- 95% of all IVT from cases covered by major press outlets have been filtered by Google

Google

INTERNAL POWER

## Next steps

- Overall AdSpam
  - Review and provide feedback on top level strategy doc (what's missing?) (*due Aug 14*)
  - Brainstorm and align on L2 strategic intents (*draft due July 10; final due Aug 14*)
  - Who: PIVs, TLV/TLS, PGMs, T&S leads
- Team-level (*due Aug 14*)
  - Develop team strategy
  - Define team level strategic intents
  - Outline projects/roadmaps to achieve strategic intents

Google

Proprietary + Confidential

Break ☕ 15 mins

PST: 8:00am - 8:15am

GMT: 4:00pm - 4:15pm

IST: 8:30pm - 8:45pm

Next session: Trust Graph Powered Analysis



Google



Privacy - Settings

# Trust Graph Powered Analysis | 30 mins

PST: 8:15am - 8:45am

GMT: 4:15pm - 4:45pm

IST: 8:45pm - 9:15pm


Google

### Building for the next generation of AdSpam impact

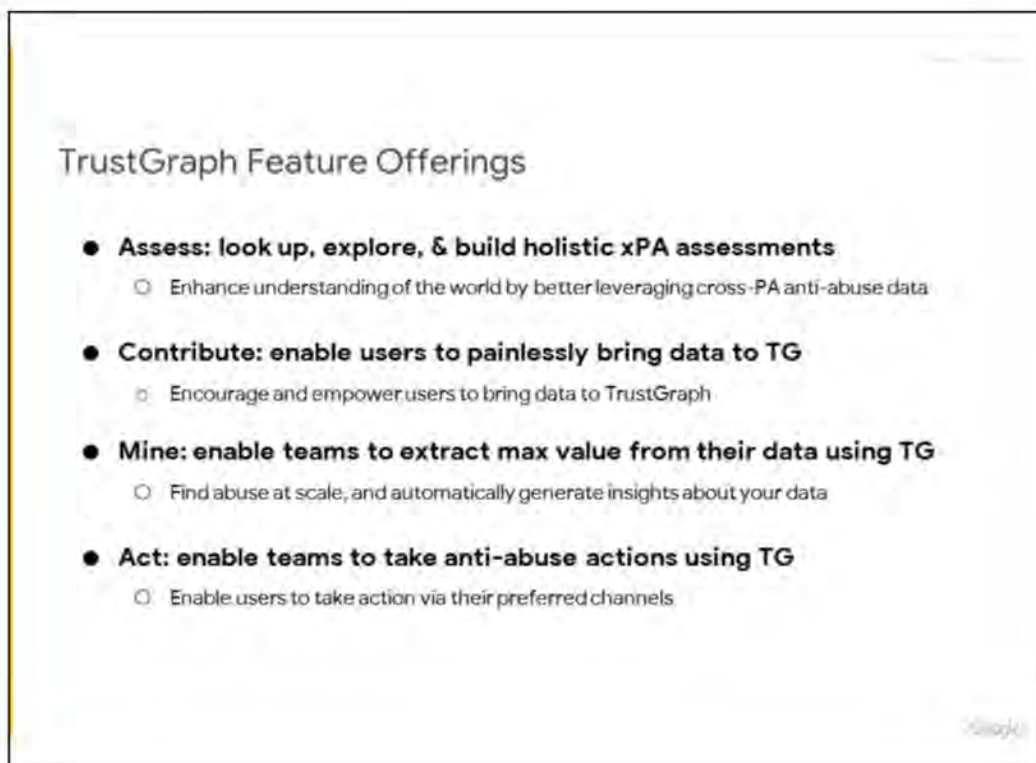
- Thesis: as our ads products continue to grow in size and complexity, and as bad actors continue to innovate, the key to our success will be our ability to scale
- There are many components to scaling (including automation), but we will always be dependent on human insights to understand the problem space
- Building tools that enable our analysts and engineers to achieve maximal impact will be a critical part of AdSpam's success
- TrustGraph, and apps using TG data (eg. TG Insights, Telescope, and SmartLeads), will be key drivers of innovation in detection and our ability to scale

Speaker: Zack


- a) educate (feedback I keep hearing is folks dont know the vision/plan)
- b) build enthusiasm (some selling on the future potential - but not overselling)
- c) feedback / alignment



Speaker: Zack



Speaker: Zack



Talking points:

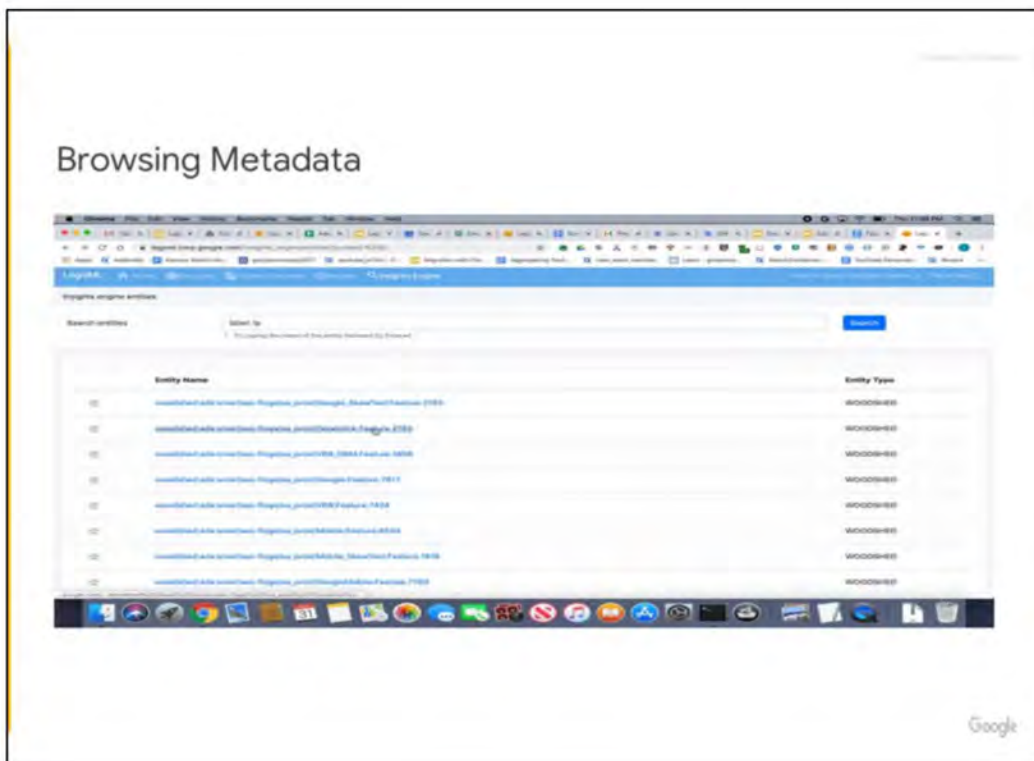
analysts can browse data using a combination of graphs and tables

mention that we have graph extraction (aka graph query language 0.5) as input  
first version of this will be Smartmeter

would like to get feedback on the possibility of integrating this with other investigation tools

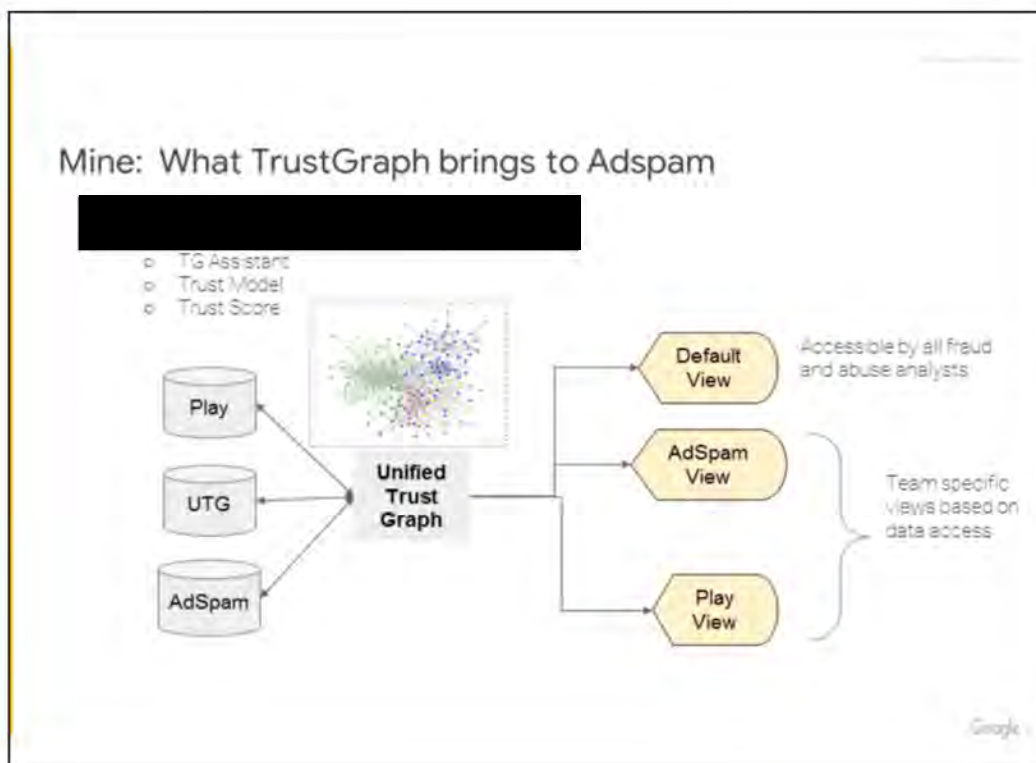
Id	Date	Text
1	06/22/2020 21:14:25	@xinzhaogoogle.com Please take a look at this slide and the next one. I plan to verbally say that these 2 UI would be combined at some point in the future in Smartmeter. Does this sound good to you?  @zacharylf@google.com for a FYI _Reassigned to Xin Zhao_
1	06/22/2020 21:14:25	SG. This is indeed our plan. The graph help visualize the most relevant clusters/nodes, while tables provide detailed info and highlights from users' interpretation on signals.



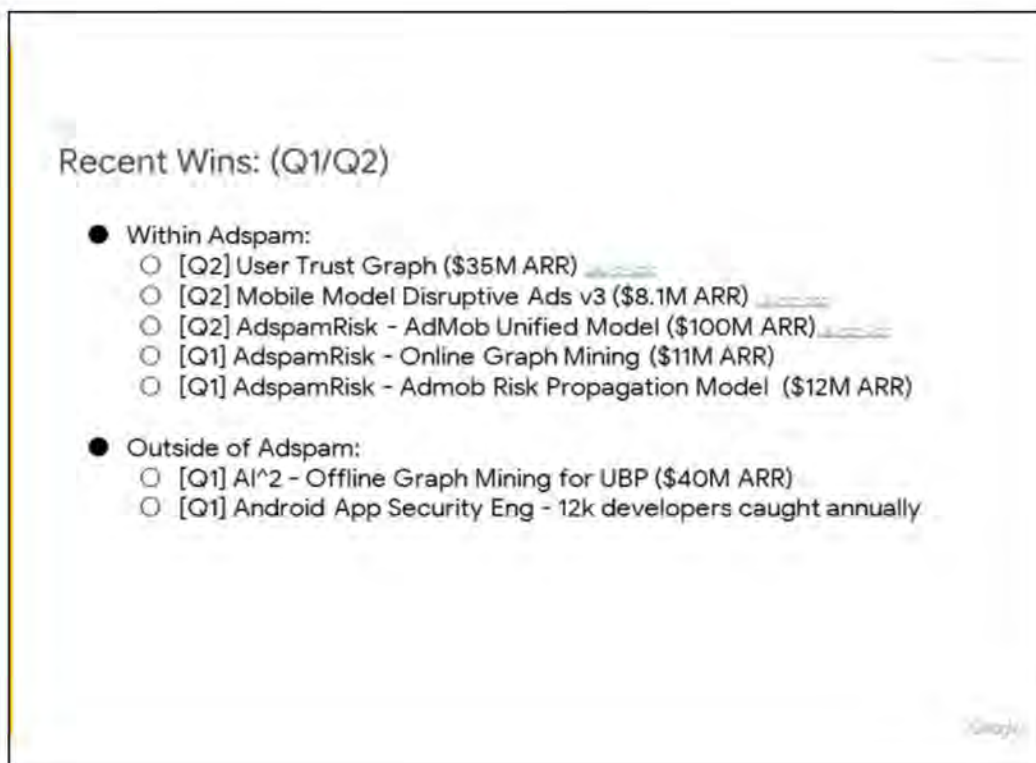


Might skip this if we don't have time





Unified Trust Graph -> Trust Model -> Trust Score

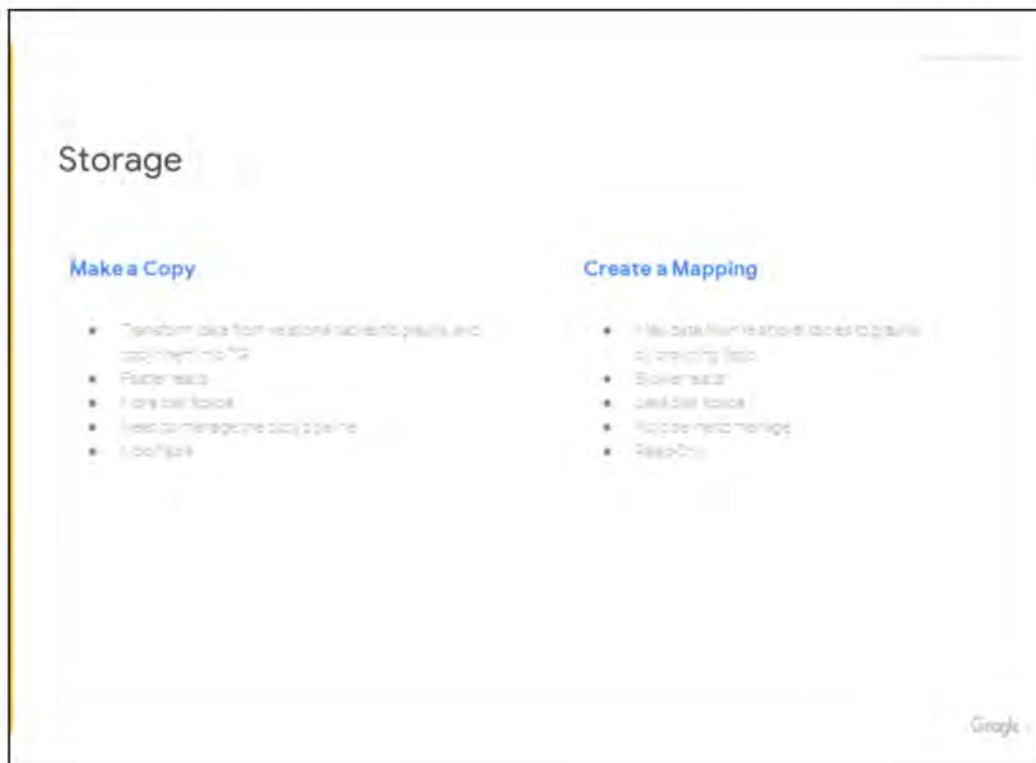


There are some launches here that we could include.

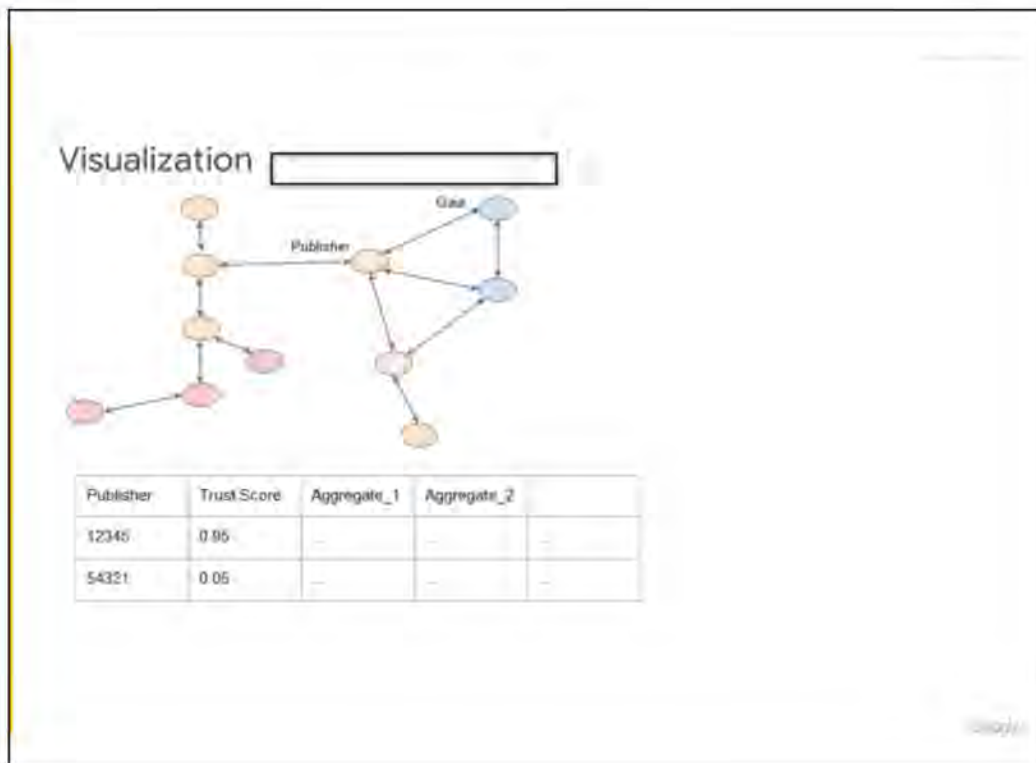
Id	Date	Text
2	06/22/2020 18:10:38	@kalpa@google.com
3	06/22/2020 21:17:47	Just added ones from here:  <a href="https://docs.google.com/document/d/1AN7Ht8j6rwtGVUYwZe0la2HI8fCbsSzF20NWHdcqs2Y/edit#">https://docs.google.com/document/d/1AN7Ht8j6rwtGVUYwZe0la2HI8fCbsSzF20NWHdcqs2Y/edit#</a>
1	06/22/2020 21:26:21	I am unable to edit but we should add these launches too - <a href="https://docs.google.com/presentation/d/1z7iyGGJEL8Kf9Rb1lly7VwTbgUH1OH568SlwM0z6BTc/edit#slide=id.g7e4601850f_1_0">https://docs.google.com/presentation/d/1z7iyGGJEL8Kf9Rb1lly7VwTbgUH1OH568SlwM0z6BTc/edit#slide=id.g7e4601850f_1_0</a>
1	06/22/2020 23:34:41	+hsiaosu@google.com Do you have any material on early wins for UTG and AdSpam Risk we can highlight here?
4	06/22/2020 23:34:41	I filled out a few more from here:  <a href="https://go/tg-impact&amp;engagement-tracker">go/tg-impact&amp;engagement-tracker</a>



Simplify this slide, stress that we'll make privacy easier.



Skip this



Talking points:

analysts can browse data using a combination of graphs and tables

mention that we have graph extraction (aka graph query language 0.5) as input  
first version of this will be Smartmeter

would like to get feedback on the possibility of integrating this with other investigation  
tools

Proprietary + Confidential

Break ☕ 15 mins

PST: 8:45am - 9:00am

GMT: 4:45pm - 5:00pm

IST: 9:15pm - 9:30pm

Next Session: Emerging Trends



Google

Privacy - Settings

# Emerging trends | 30 mins

PST: 9:00am - 9:30am  
GMT: 5:00pm - 5:30pm  
IST: 9:30pm - 10:00pm

Google





The goal of this session is to seed a conversation about the big picture trends and challenges facing AdSpam as we go about our 2021 planning.

This isn't a TED talk, I don't have a crystal ball, but hopefully this will start a useful discussion and help us frame our annual strategies in a way to address some of these challenges.

AdSpam operates in a complex and dynamic space

- The Digital Ads Ecosystem is constantly growing and evolving
- Google is itself a large, complex organization with a wide array of products both new and established needing AdSpam support.
- AdSpam has been relatively successful in standardizing our policies and requirements to enable us to scale across Google's ads business and integrate most new products.
- However, there are an increasing number of new products and requests that don't fit neatly into our standardized solution, requiring custom work and constant realignment.

**The main drivers of complexity appear to be accelerating.  
What else can we do to get ahead of the problem?**

AdSpam operates in a really complex space. Externally, the digital ads ecosystem is growing, evolving, and maturing - reaching new forms of media.

Internally, Google's variety of businesses, infrastructure, and organizational structure is extremely complex in its own right.

We've developed ways to help us navigate complexity (like alignment), and address it more structurally through standardizing our integration requirements, and innovating in our analysis methodologies and infrastructure (eg, ASML, TrustGraph, Organic Labeling).

This has helped us be relatively successful in both supporting new product requirements and feature requests, and scaling our defenses across Google's wide array of businesses, as part of our "standard solution"

The emerging trends that are driving the complexity we deal with on a daily basis, originating both internally and externally to Google, are pointing to an expansion in variety and diversity of online ad experiences for users and platforms where they happen.

As a result, we're increasingly seeing requests fall outside of these standard processes and solutions, requiring custom work and constant alignment.

As we go through strategic planning, how can we get ahead of these trends? Are there any asymmetric bets we can invest in in 2021?

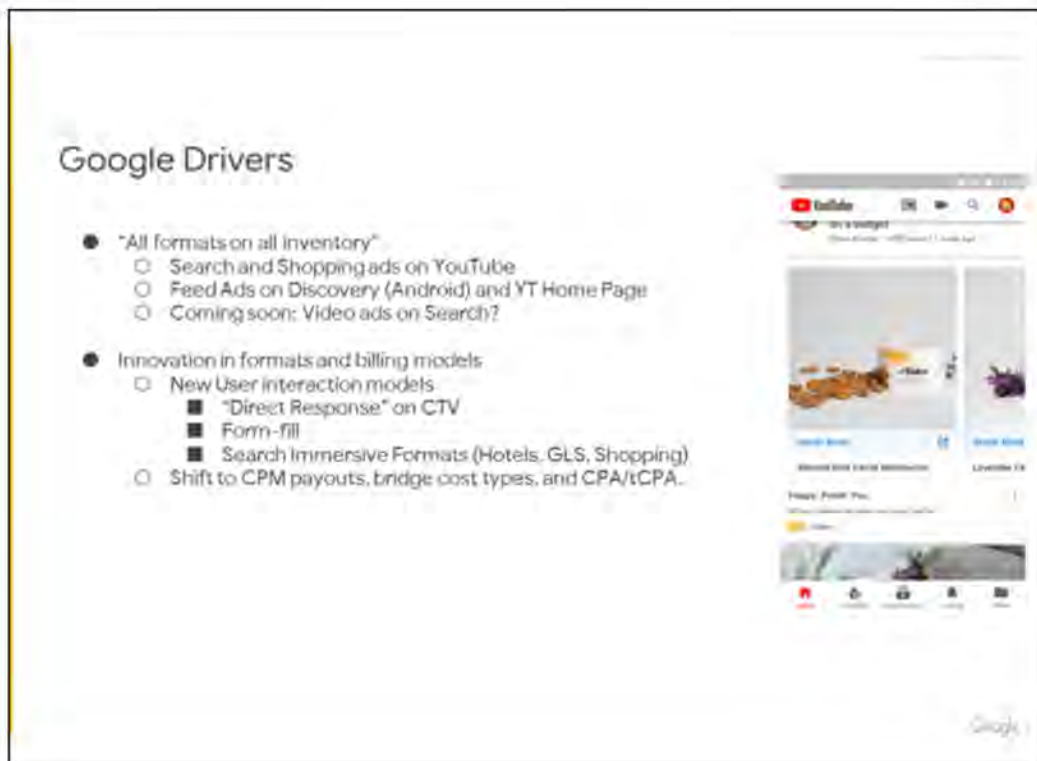


Starting with external trends:

The ads ecosystem follows user behavior. Increasingly, users are engaging with ad supported content across a higher number and wider variety of devices. We've all heard about the shift from desktop to mobile, and now we're seeing the beginnings of a shift to wearables, smart speakers, and connected TV devices. New devices, ad formats, attack vectors, signals and collection mechanisms, all require additional attention from AdSpam.

The second trend to call out here is similar - as the digital ads ecosystem matures, we're starting to see a shift in "traditional" ad spend moving to digital, and from digital to programmatic. While CTV, DOOH, and Audio are all multi-billion-dollar markets already, look at how much is still spent on their traditional siblings -- over \$250B!

Not all of this will move to digital, and what does isn't going to move all at once, but I think this does at least hint that we'll continue to see continued growth in these new areas for the years to come.




Within Google, we're increasingly seeing an effort to consolidate ad format and inventory offerings, to allow advertisers to run the same ad in more places. For example, if you look at recent developments in the YouTube app, you'll notice AFS ads, Shopping ads, and Feed ads have been incorporated into the Home and Search pages.

At the same, innovation in formats and billing models are challenging some of AdSpam's assumptions around expected user interactions, metrics, and even logs and event-hierarchies.

For example, in immersive formats on Search such as Shopping, Local Services, or Maps, we often log multiple clicks for each ad - the first is to open the immersive, the second is the ad click. Enhanced Conversions have found that their requirements can't be met with the AdEvents logs.





By way of example: CTV, Audio, and DOOH are the “next big thing”, and incorporate a lot of these trends.  
[DEFINE THE TERMS]

We’re the furthest along in our defense of CTV, and so far, we’ve tried to apply the standard AdSpam playbook.

As we expand down the long-tail of CTV devices, and into smart speakers and billboards, do we need a different approach? How does the definition of IVT change for these formats?

Recent shifts in Mobile app advertising and related abuse vectors are also rapidly changing, and challenging us in unique ways - I believe Zack is going to cover some of these in a future session.



So how do these trends affect AdSpam?

These trends are pointing to two new business requirements that we'll need to find a way to support.

## AdSpam Challenges

These new requirements will pose some significant challenges for AdSpam:

- **Signal Collection:** As the number and variety of device platforms and implementations needed to support the ads business expands, it will be difficult to ensure adequate and consistent signal collection.
- **Non-Standard Implementations:** As more products require custom work from AdSpam, it will increasingly be a challenge to scale and maintain our defenses across all of Google's businesses.
- **Organizational Alignment:** As formats, signal collection mechanisms, and device platforms shift, does our organizational structure and vertical/horizontal team structure set us up for success?
- **Elevated Operational Costs:** Any gaps or shortcomings in our approach to defending new products increases the operational burden on our T&S teams.

Google

And becoming flexible enough to support those new requirements won't be easy.



## Key Questions for Strategic Planning

- How do we build for flexibility, while still supporting new products?
- How do we maintain defense quality in signal-poor environments?
- How do we impact product roadmaps, and/or decide which products to support?
- What is the right balance to strike between defending a product area where it is, vs. requiring it to change before we can offer support?
- Can we come up with a new engagement model that enables us to scalably support more of the requests currently considered "custom"?
- Are there any asymmetric bets we can invest in today to help us prepare to meet these challenges?

So as we go through this strategic planning summit and related exercises, I think there are a few key questions that I hope we can keep in the back of our minds.

Proprietary + Confidential

# Improving precision | 30 mins

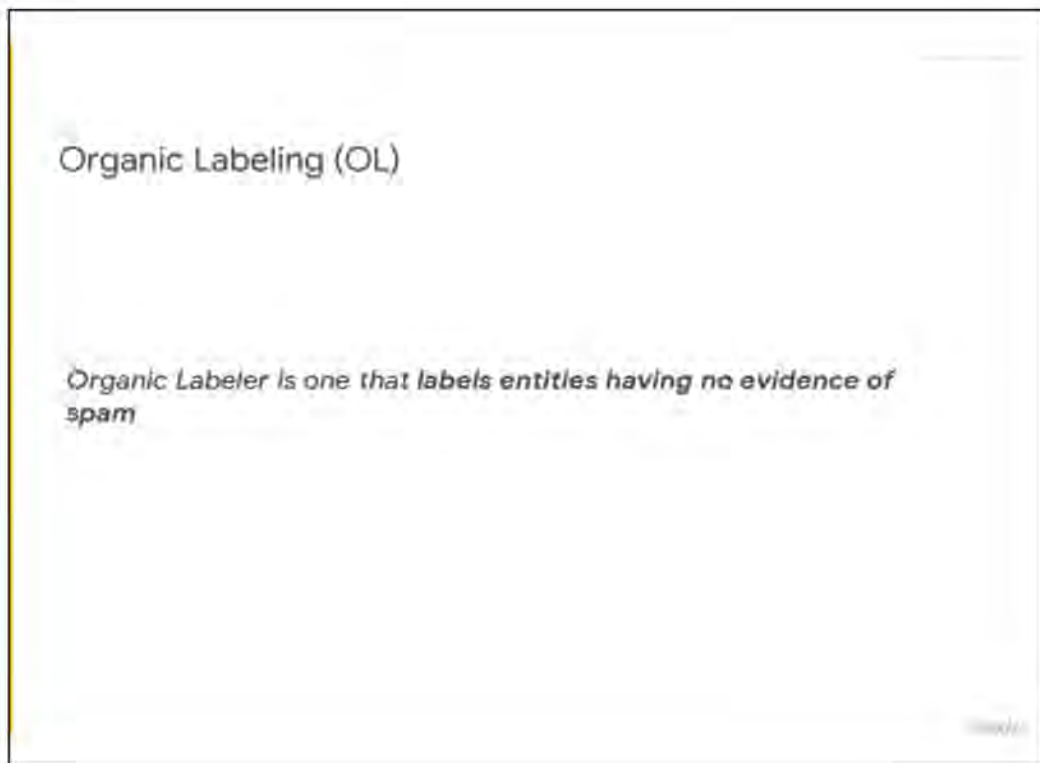
PST: 9:30am - 10:00am

GMT: 5:30pm - 6:00pm

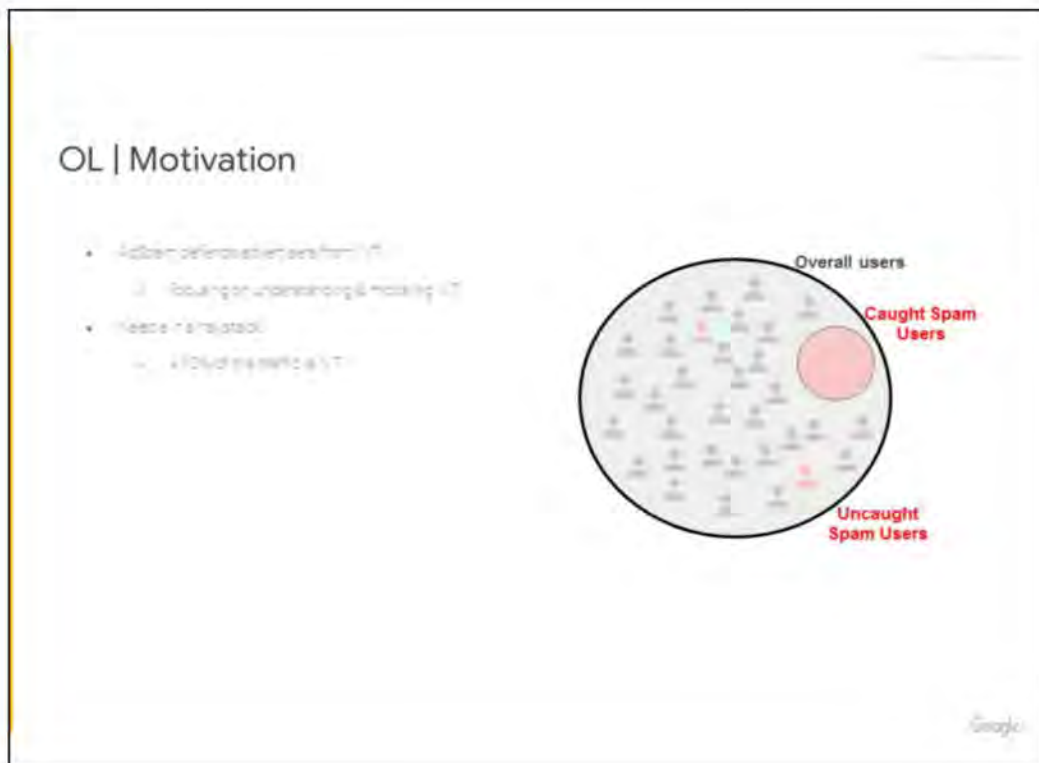
IST: 10:00pm - 10:30pm

Google

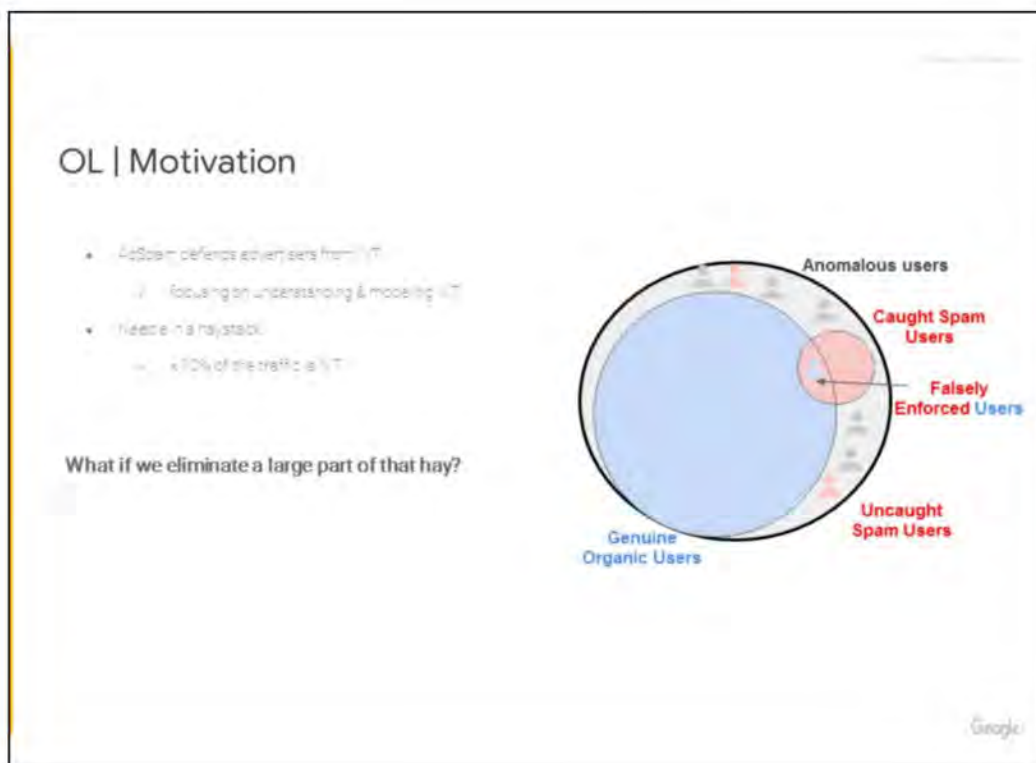
Id	Date	Text
1	06/16/2020 22:03:35	@spacemanspiff@google.com @atali@google.com <small>Proprietary + Confidential</small>
1	06/16/2020 22:03:35	@sulabha@google.com I think you had a couple slides to share too



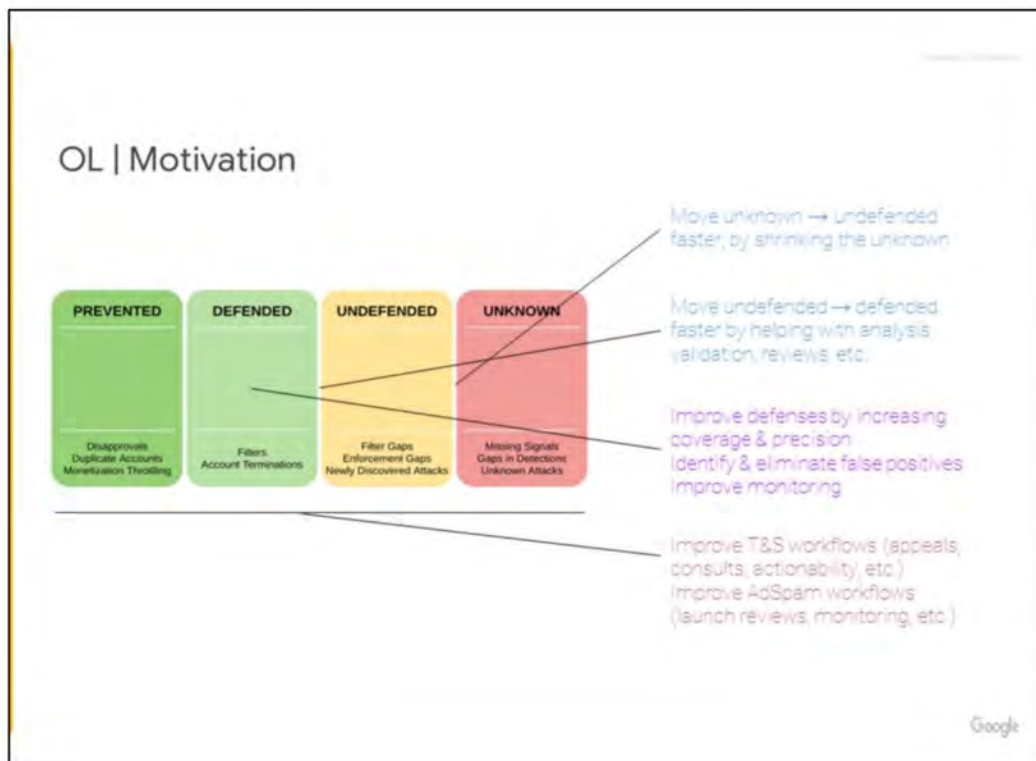
Slides heavily stolen from a bunch of different slide decks by rahulmah@.



Uncaught spam << 0.5%.

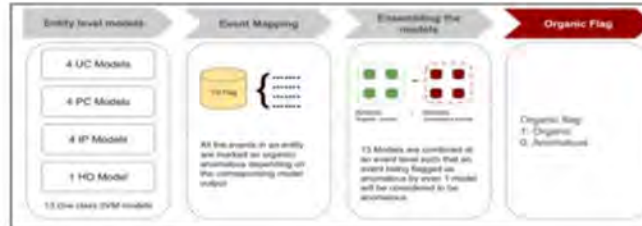


Uncaught spam << 0.5%.



## OL | s-spam

- Launched in Q1 2020
  - One-Class Support Vector Machines (OCSVM) / ML-sklearn
  - Unsupervised (no ground truth)
  - Mimics how we look for organicness in manual reviews
    - Learns the majority class function (majority of our traffic is organic), and predicts what is similar to the majority class
  - Labels available in AdSense and GoldStone



Google







rahulmah: As per OL insights we know that at least 30-40% FPs are in form of collateral damage and turn downs are not the complete solution? We expect this to be the case across verticals as well.

## OL | There's more

- Identification of time-consuming organic traffic (SDDT, TR, etc.)
  - Identifying automatic disavows
- Improve DryBar's performance
- Optimization of many T&S processes (eg. prioritization by actionability)
- Insights into which buckets of traffic we need to focus on
- Monitoring ASOR, UOR
- Streamlining manual reviews & auction reviews
- Integration with Case Trust Graph

Google

## OL | Takeaways

What we hope you'll remember this week:

- Organic Learning is *green*
  - promising solution to many important problems of AdSense
- We need to keep going
  - utilize the know-how that took a long time to build
  - utilize the infra and invest into its scalability
  - investment should start in H2 2025, with all verified having OL impact in 2026
- The full impact of OL won't be realized until
  - all verified have OL ad that
  - OL becomes a *core muscle* of AdSense

Google

Proprietary + Confidential

Break ☕ 15 mins

PST: 10:00am - 10:15am

GMT: 6:00pm - 6:15pm

IST: 10:30pm - 10:45pm

Next Session: Impression Spam



Google

Proprietary + Confidential

# Impression Spam | 30 mins

PST: 10:15am - 10:45am

GMT: 6:15pm - 6:45pm

IST: 10:45pm - 11:15pm

Google

## Impression Spam - Problem Statement

*Ad impressions are much less well-defended across Google's ad products than are ad clicks...*

*...yet impressions represent > 50% of total monetized ad spend.*

*As of early 2020, the dollar-weighted pure CPM-based IVT rate is 1.6%,  
vs CPC-based IVT rate of 8.1%*

*Meanwhile, publishers already have, and face new incentives, to inflate impressions:  
Impressions are already the event by which GAM publishers are paid for  
monetization, and will be for AdMob in 2020, and AdSense in 2021.*

*Google*

Speaker: Zack


Impression Spam - But How Big?

**Extra Sensitive!**

We estimate:

- Upto \$300M ARR uncaught IVT on pure CPM cost-type*
- Upto \$470M ARR in net loss to Google on the CPC-to-CPM bridge*

*Independent analyses* *Supported by multiple,*



Speaker: Zack



## Impression Spam - Vision and Goals

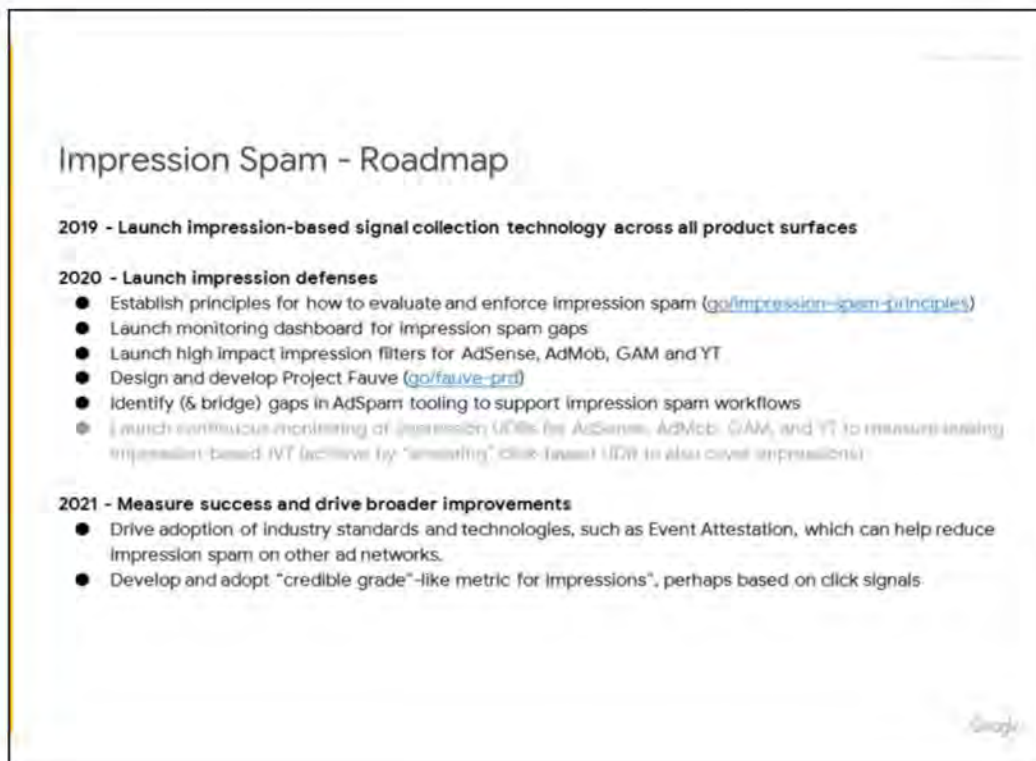
**Vision:** *IVT on ad impression is as equally well defended as IVT on ad clicks.*

Goals, in order of priority:

- 1) We do not charge advertisers for invalid impressions, wasting spend and eroding ROI.
- 2) We do not pay publishers for invalid impressions, incentivizing more fraud and abuse.
- 3) We protect both advertisers from underfiltration and publishers from overfiltration.

Google

Speaker: Zack



Speaker: Arun

## Impression Spam 2020

- **Formed Impression Spam Working Group** in Eng. and in T&S
- **WhiteOps Integration based augmented defenses**
  - Albus 3PE (Q3)
- **Direct Impression defenses** (aka native defenses)
  - GDA pCookie/pCTR models doubled query spam rate from 1.08% to 2.00% ([verdict\(189312\)](#))
  - Mobile pCookie/pCTR models (Q3)
  - GVP pCookie/pCTR (H2, maybe)
- **Transfer Learning based Impression defenses**
  - GDA Click to Query impression defenses (Q3)
  - Mobile Click to Query (H2)
- *Imputations*
- *Preempt*

**Call to Action:**

- **Understand** how CPM and CPC-to-CPM impacts and affects your products
- **Help us** improve our measurement of impression spam, and understanding of top attack vectors
- **Drive** impression spam defenses on your networks, leveraging product specific signals, and shrink the gap

Google

Speaker: Arun?

## Challenges with reviewing Impression Traffic

- Metrics** Lack of strong validation/re-tuning traffic metrics like CG (like credible grade for clicks)
  - Increased review time
  - Risk of missing significant clusters of spammy traffic.
- Signals** Increased dependencies on signals outside of traditional logs (Chrome-ATM, GodCluster)
  - Overfiltration risk due to non traffic signals used as spam indicators
- Tools** Query information split in multiple log sources which slows down investigations
- Tools** Most of the infra has limited support for Queries.
  - e.g., Woodshed has 10% of 1% AEQ data which can be misleading
- Tools** Tools like Dojo & WPI provide limited support for query or impression level action
- Research** Limited coverage, understanding & research for existing impression signals
- Research** Understanding of impression spam not on par with click spam across the team, missing case studies and similar documentation

Speaker: Prachi

## Impression Spam Principles - A Re-Think

- 1) **Detecting Invalid Impressions at Scale Requires Using Downstream Events**
  - a) TLDR: Can we create an "ImpressionCredibleGrade" to assess baskets of events?
- 2) **The Law of Spammy Entities - Part 1 - Verdict Transferability**
  - a) TLDR: If the clicks associated with a particular entity are spammy, our default position should be to assume the impressions are invalid
- 3) **The Law of Spammy Entities - Part 2 - Pipeline Transferability**
  - a) TLDR: If you determine a particular entity has invalid clicks and impressions on GDA, you should bias towards blacklisting this entity on other demand sources (eg. DV360)
- 1) **Impression Spam Requires an Appropriate Precision Bar**
  - a) TLDR: We should determine an acceptable precision bar for impression spam defenses, which may be lower than for click defenses

[go/impression-spam-principles](https://www.google.com/go/impression-spam-principles)

Google

Speaker: Zack

## Impression Spam Solutions - Les Fauves!

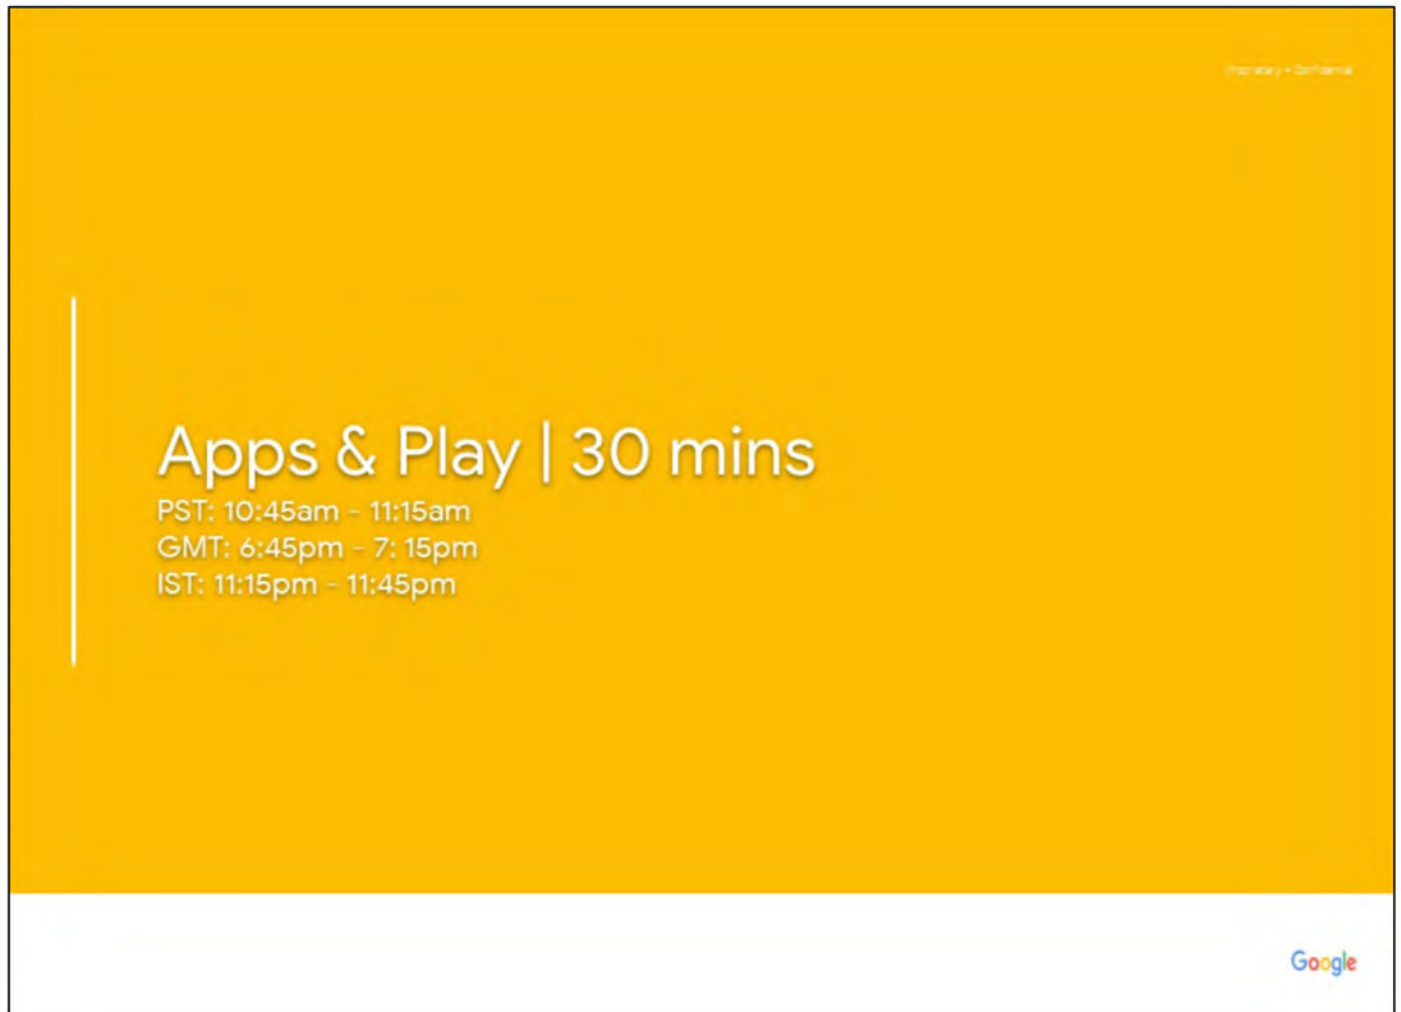
- While we aim to launch impression-event classification, there may be limits
  - Transfer learning, imputation, entity labeling - all will help significantly
  - Impression metrics will also help
  - BUT there are probably limits that will prevent complete parity with clicks
- But we can also use auction levers to a) better defend and b) gain information
  - Eg. Project Fisher prevented conv spam, but not by better marking conversion spam
- Proposed solution - Project Fauve (what's in a [name](#) anyway?)
  - Develop risk models that identify publishers likely to drive impression spam
  - Apply mechanism to shift demand pre-auction from CPM to CPC
  - Metrics, models, and dashboards to interpret observed changes post-treatment

**TLDR: instead of better classifying invalid impression events, lets classify risky publishers, and require them to generate clicks in order to generate revenue**

[go/fauve-prd](#)

Google

Speaker: Zack





Speaker: Zack



ASE + AdSpam Joint-2020 AdFraud Roadmap

**2020 AdSpam+Apps Safety PA-level OKR**  
Reduce ad fraud on Play by \$100M ([go/adfraud2020](#))

1. Reduce latency & expand enforcement
  - a. Align enforcement principles and practices across top abuse sectors
2. Programmatic signals, verdicts, & leads sharing
  - a. Eliminate organizational information gaps and maximally leverage signals and leads across-PA's
3. Improve detection & classification
  - a. Dramatically improve both independent and joint ad fraud detection & classification recall

*Google*

Speaker: Zack

2020 Jan 1 State

	Reduce latency & expand enforcement	Programmatic signals, verdicts, & leads sharing	Improve detection & classification
Summary	Align enforcement <b>principles and practices</b> across top abuse vectors	<b>Eliminate organizational information gaps and maximally leverage</b> signals and leads from across PAs	Dramatically <b>improve both independent and joint ad fraud detection &amp; classification</b> recall
2019 EOY State	<b>State: weak</b> Ad-hoc alignment on large enforcement cases. Largely lacking principle and process alignment.	<b>State: improving</b> Bi-directional sharing largely offline, via spreadsheet. Some investment with Trust Graph onboarding and leads sharing with SmartLeads.	<b>State: improving</b> Historically weak on top abuse vectors. Large launches and momentum building through 2019. On track for large 2020 gains.
Target 2020 EOY State	Aligned on shared principles with enforcement processes deployed across teams	Enforcement status, core abuse signals, risk scores, classification outputs, and prioritized leads from ASE, AdSpam, and T&S universally available in Trust Graph and SmartLeads	Individually and jointly-owned detection and classification techniques launched, providing coverage for > x% of enforced ad fraud on Play

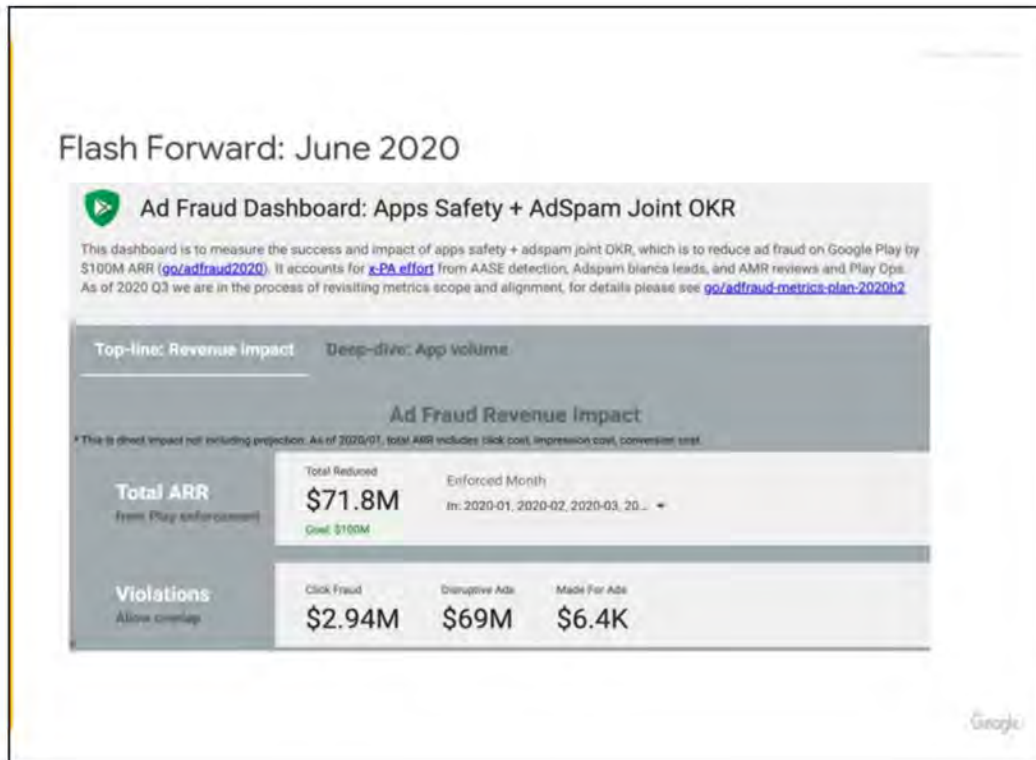
Speaker: Zack

ASE + AdSpam Joint-2020 AdFraud Roadmap			
	Reduce latency & expand enforcement	Programmatic signals, verdicts, & leads sharing	Improve detection & classification
Key Projects	<ul style="list-style-type: none"> <li>Launch joint-ad fraud enforcement principles</li> </ul>	<ul style="list-style-type: none"> <li>Onboard top signal sharing use-cases onto Trust Graph</li> </ul>	<ul style="list-style-type: none"> <li>Launch new signal collection technologies to improve detection</li> </ul>
	<ul style="list-style-type: none"> <li>Shared T&amp;S enforcement processes</li> </ul>	<ul style="list-style-type: none"> <li>Onboard AdSpam, ASE, and T&amp;S app/dev classifications and enforcement verdicts onto Trust Graph</li> </ul>	<ul style="list-style-type: none"> <li>Launch improved classifiers using new in-house and cross-PA signals and labels</li> </ul>
	<ul style="list-style-type: none"> <li>Launch automated Marmot ortho sooner atop SmartLeads</li> </ul>	<ul style="list-style-type: none"> <li>Launch SmartLeads to enable cross-PA leads and cluster sharing</li> </ul>	
	<ul style="list-style-type: none"> <li>Launch AdSpam pre-bid filter using ASE APK blacklists</li> </ul>	<ul style="list-style-type: none"> <li>Ad fraud focused GP rewards program to help generate more leads/intelligence</li> </ul>	
	<ul style="list-style-type: none"> <li>On-device warnings for top abuse vectors</li> </ul>		

[go/antifraud2020-roadmap](#)

Google

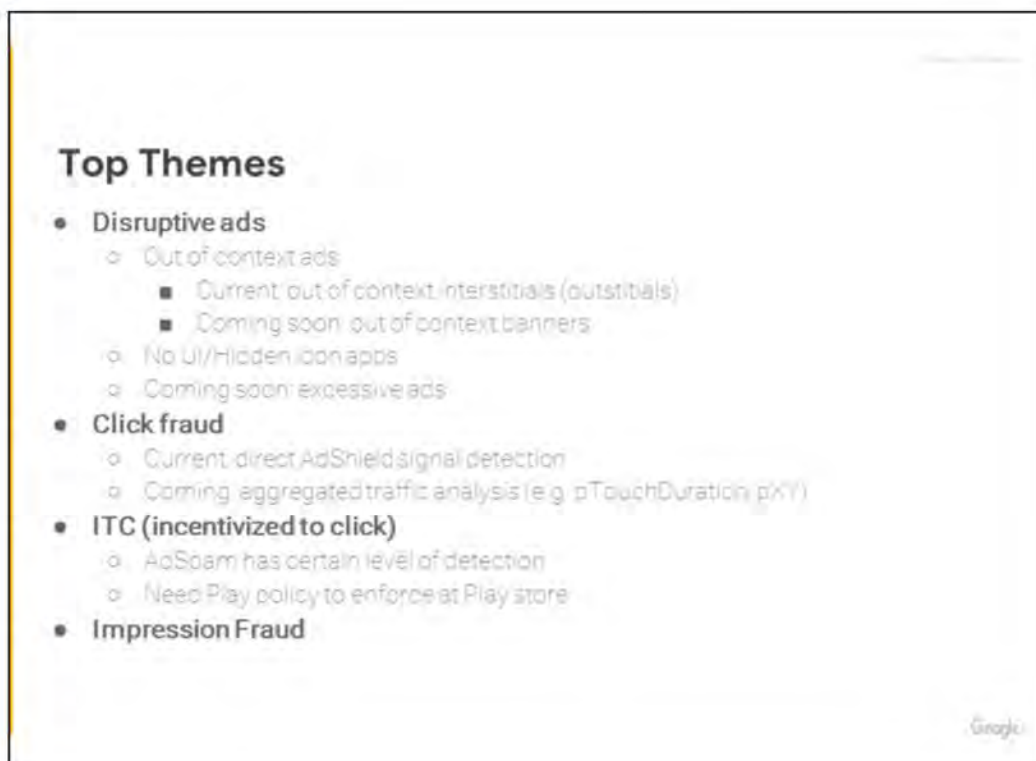
Speaker: Zack



Speaker: Zack

2020 H2 Roadmap			
	Reduce latency & expand enforcement	Programmatic signals, verdicts, & leads sharing	Improve detection & classification
Summary	Align enforcement <b>principles and practices</b> across top abuse vectors	<b>Eliminate organizational information gaps and maximally leverage</b> signals and leads from across PAs	Dramatically <b>improve both independent and joint</b> ad fraud <b>detection &amp; classification</b> recall
2020H1 Highlights	<ul style="list-style-type: none"> <li>- Aligned on new <a href="#">Ad Fraud definition</a> expected to go live in July</li> <li>- Enforcement expanded on Google Play using the <a href="#">Credible Technical Evidence</a> policy.</li> </ul>	<ul style="list-style-type: none"> <li>- <a href="#">SmartLeads</a>-Marmot integration completed</li> <li>- Play Auto-suspend launched for high precision flags from AdSpam</li> <li>- Joint enforcement against <a href="#">Cheetah Mobile</a> from AdSpam and Play for Disruptive Ads: \$45M AAR</li> </ul>	<ul style="list-style-type: none"> <li>- Launched <a href="#">AdFraud2020 Dash</a></li> <li>- Landed disruptive ads dynamic detection migration from AppCrawler to Marmoset</li> <li>- Landed <a href="#">Bianca</a> disruptive ads traffic based detection with AdMob data</li> <li>- Landed App <a href="#">CoClick</a> Filter: \$4.2M ARR</li> <li>- Landed <a href="#">Caprice</a> filter (pTouchDuration) in ramp up: \$19.7M ARR</li> <li>- Developed Unified <a href="#">PubAppRisk</a> model: \$140M ARR</li> <li>- Launched Bianca disruptive ads SmartLeads generator</li> </ul>
2020H2 Plans	<ul style="list-style-type: none"> <li>- El Paso policy bundle (Oct 25) and further enforcement practice alignment.</li> <li>- Inactive ads policy roll out and enforcement implementation with Detox</li> </ul>	<ul style="list-style-type: none"> <li>- <a href="#">Magellan</a> project: improving auto detection and using AdSpam signals to improve Marmot ability to detect</li> <li>- Overhaul of how ASE APK categorization are applied on ads serving</li> </ul>	<ul style="list-style-type: none"> <li>- Developing impression defences (expecting &gt;\$50M ARR)</li> <li>- Multisource attribution and measurement in Ad Fraud tracking</li> <li>- Additional <a href="#">SmartLeads</a> ad fraud generators</li> </ul>
2020EOY	ON TRACK	ON TRACK	ON TRACK

Speaker: Zack



Speaker: Xin

## AdSpam, Play, Android Collaboration

- **Leadgen engine:** generate entity-level leads from our event-level detection
- **SmartLeads:** end-to-end support for leads sharing
  - o leads signal compilation, sharing, result tracking, dashboard
- **Magellan:** an intelligent pipeline to run apps and detect ad fraud activities
  - o Use ads traffic info to guide app automation
  - o Use app automation to save lead/validation cost
  - o Collect ad fraud evidence efficiently
- **Bi-directional feedback loop**
  - o Use AdSpam leads to help improve Play/detection
  - o Use Play/detection feedback to improve AdSpam/leadgen rules

Google

Speaker: Xin

## Challenges in Current Collaboration

- **Ad fraud mechanisms become more sophisticated**
  - Highly packed code
  - Cloaking
  - Blend organic and spam traffic
  - Real humans to commit ad fraud instead of pure program
- **AdSpam TP, but not Play actionable**
  - E.g. VM-based or crowdsourcing attack
- **Policy not aligned or missing**
  - E.g. disruptive ads definition, no ITC policy

Google

Speaker: Xin

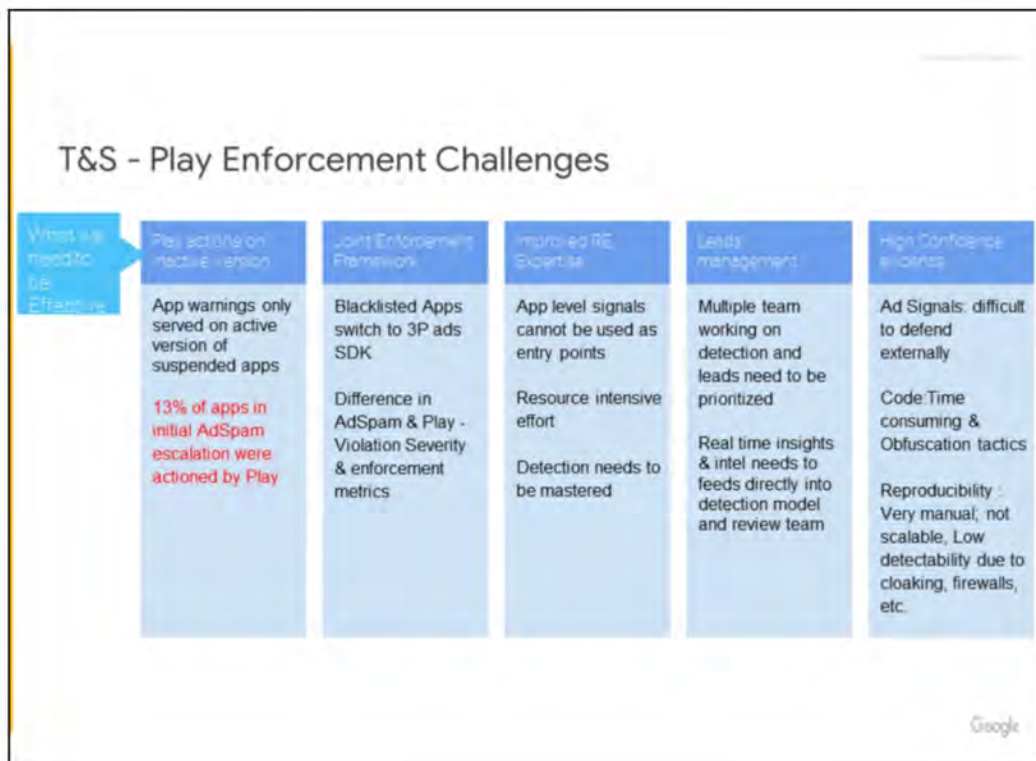


## Future Tasks

- **Handling more sophisticated ad fraud**
  - Highly packed code
  - Cloaking
  - Blend organic and spam traffic
  - Real humans to commit ad fraud instead of pure program
- **Improve signals accuracy**
  - Highly complex and diversified ad products
  - Signals not well polished
- **Align and augment Play policies**
  - E.g. different disruptive ads definition, no ITC policy
- **Respond to threats from new areas**
  - Impression fraud, Webview based attack

Google

Speaker: Xin



Speaker: Prachi

Proprietary + Confidential


# Closing remarks | 15 mins

PST: 11:15am - 11:30am

GMT: 7:15pm - 7:30pm

IST: 11:45pm - 12:00am

Google



Reminder to provide feedback



Reminder to provide feedback